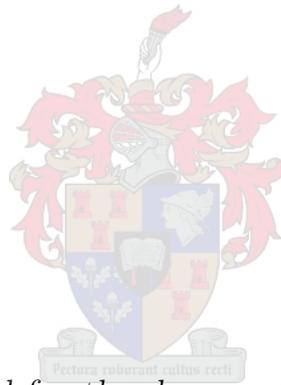


Virtual Reality Assisted Fluorescence Microscopy Data Visualisation and Analysis for Improved Understanding of Molecular Structures Implicated in Neurodegenerative Diseases

by

Rensu Petrus Theart



*Dissertation presented for the degree of Doctor of Philosophy
of Engineering (Electronic) in the Faculty of Engineering at
Stellenbosch University*

Supervisor: Prof. TR. Niesler

Co-supervisor: Prof. B. Loos

March 2020

Declaration

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: March 2020

Copyright © 2020 Stellenbosch University
All rights reserved.

Abstract

Virtual Reality Assisted Fluorescence Microscopy Data Visualisation and Analysis for Improved Understanding of Molecular Structures Implicated in Neurodegenerative Diseases

R.P. Theart

*Department of Electrical and Electronic Engineering,
University of Stellenbosch,
Private Bag X1, Matieland 7602, South Africa.*

Dissertation: PhD (Electronic)

March 2020

Confocal microscopy is one of the major imaging tools used in molecular life sciences. It delivers detailed three-dimensional data sets and is instrumental in biological analysis and research where structures of interest are labelled using fluorescent probes. Usually, this three-dimensional data is rendered as a projection onto a two-dimensional display. This can however lead to ambiguity in the visual interpretation of the structures of interest in the sample. Furthermore, analysis and region of interest (ROI) selection are also most commonly performed two-dimensionally. This may inadvertently lead to either the exclusion of relevant or the inclusion of irrelevant data points, consequently affecting the accuracy of the analysis.

We present a virtual reality (VR) based system that allows firstly, precision region of interest selection and colocalisation analysis, secondly, the spatial visualisation of the correlation of colocalised fluorescence channels and, thirdly, an analysis tool to automatically determine the localisation and presence of mitochondrial fission, fusion and depolarisation.

The VR system allows the three-dimensional reconstructed sample data set to be interrogated and analysed in a highly controlled and precise manner, using either fully-immersive hand-tracking or a conventional handheld controller. We apply this system to the specific task of colocalisation analysis, an important tool in fluorescence microscopy. We evaluate our system interface by means of a set of user trials and show that, despite inaccuracies associated with the hand tracking, it is the most productive and intuitive interface compared to the handheld controller.

Applying the VR system to biological sample analysis, we subsequently calculate several key colocalisation metrics using both two-dimensionally and three-dimensionally derived super-resolved structured illumination-based data sets. Using a neuronal injury model, we investigate the change in colocalisation between two proteins of interest, Tau and acetylated α -tubulin, under control conditions as well as after 6 hours and again after 24 hours of neuronal injury. Applying the VR based system, we demonstrate the ability to perform precise ROI selections of 3D structures for subsequent colocalisation analysis. We demonstrate that performing colocalisation analysis in three dimensions enhances its

sensitivity, leading to a greater number of statistically significant differences than could be established when using two-dimensionally based methods.

Next, we propose a novel biological visual analysis method for the qualitative analysis of colocalisation. This method visualises the spatial distribution of the correlation between the underlying fluorescence channel intensities by using a colourmap. This method is evaluated using both synthetic data and biological fluorescence micrographs, demonstrating enhancement of the visualisation in a robust manner by indicating only truly colocalised regions.

Mitochondrial fission, fusion and depolarisation events are important in cellular function and viability. However, the quantitative analysis linked to the localisation of each event in the three-dimensional context has not been accomplished. We extend the VR system to analyse fluorescence-based time-lapse sequences of mitochondria and propose a new method to automatically determine the location and quantity of the mitochondrial events. The detected mitochondrial event locations can then be superimposed on the fluorescence *z*-stacks. We apply this method both to control samples as well as cells that were treated with hydroxychloroquine sulphate (HCQ) and demonstrate how a subsequent quantitative description of the fission/fusion equilibrium as well as the extent of depolarisation can be determined.

We conclude that virtual reality offers an attractive and powerful means to extend fluorescence-based microscopy sample navigation, visualisation and analysis. Three-dimensional VR-assisted ROI selection enable samples to be interrogated and assessed with greater precision, thereby exploiting the potential of fluorescence-based image analysis, such as colocalisation, in biomedical research. The automatic localisation and quantification of mitochondrial events can support research of mitochondrial function in healthy and diseased cells, where quantitative analysis of fission, fusion and depolarisation is of importance.

Uittreksel

Virtuele Realiteit Aangehelpde Fluoresensie-Mikroskopie Datavisualisering en Analise vir Verbeterde Insig in Molekulêre Strukture Betrek in Neurodegeneratiewe Siektes

(“Virtual Reality Assisted Fluorescence Microscopy Data Visualisation and Analysis for Improved Understanding of Molecular Structures Implicated in Neurodegenerative Diseases”)

R.P. Theart

*Departement Elektries en Elektroniese Ingenieurswese,
Universiteit van Stellenbosch,
Privaatsak X1, Matieland 7602, Suid Afrika.*

Proefskrif: PhD (Elektronies)

Maart 2020

Konfokale mikroskopie is een van die belangrikste beeldinstrumente wat in die molekulêre lewenswetenskappe gebruik word. Dit lewer gedetailleerde drie-dimensionele datastelle wat instrumenteel is vir biologiese analise en navorsing waar strukture van belang gemerk word met behulp van fluoresserende probes. Gewoonlik word hierdie drie-dimensionele data as 'n projeksie op 'n twee-dimensionele skerm weergegee. Dit kan egter lei tot dubbelsinnigheid in die visuele interpretasie van die struktuur van belang in die biologiese monster. Verder word analise en streek van belang (SVB) seleksie ook meestal twee-dimensioneel uitgevoer. Dit kan onbedoeld lei tot die uitsluiting van relevante of die insluiting van irrelevante datapunte, wat die akkuraatheid van die analise kan beïnvloed.

Ons bied 'n virtuele realiteit (VR)-gebaseerde stelsel aan wat eerstens presisie SVB seleksie en kolokaliseringanalise in staat stel, tweedens, die ruimtelike visualisering weergee van die korrelasie van gekolokaliseerde fluoressensie-kanale, en derdens 'n analise-metode vir die outomaties lokaliserings van mitochondriale splitsing, samesmelting en depolarisasie.

Die VR-stelsel laat toe dat die drie-dimensionele gerekonstrueerde monsterdatastel op 'n hoogs gekontroleerde en akkurate wyse ondersoek en geanaliseer kan word deur gebruik te maak van 'n volledig immersiewe hand-volg toestel óf 'n konvensionele spelbeheerder. Ons pas hierdie stelsel toe op die spesifieke taak van kolokaliseringanalise, 'n belangrike instrument in fluoressensie-mikroskopie. Ons evalueer die twee stelsel-koppelvlakke aan die hand van 'n stel gebruikersproewe en wys dat die hand-volg toestel, ondanks onakkuraatheid wat met die dit verband hou, die mees produktiewe en intuïtiewe koppelvlak is in vergelyking met die spelbeheerder.

Met die toepassing van die VR-stelsel op biologiese monsteranalise, bereken ons daarna verskeie sleutel kolokalisasiemates met behulp van beide twee-dimensionele en drie-dimensionele “super-resolved structured illumination” gebaseerde datastelle. Met behulp van 'n

neuronale beseringsmodel ondersoek ons die verandering in kolokalisasie tussen twee proteïene van belang, Tau en geasetileerde α -tubulien, onder beheerstoestande sowel as na 6 uur en weer na 24 uur na neuronale besering. Met gebruik van die VR-gebaseerde stelsel, demonstreer ons die vermoë om presiese SVB-seleksies van 3D-strukture uit te voer vir latere kolokaliseringsanalise. Ons demonstreer dat die uitvoering van kolokaliseringsanalise in drie dimensies die sensitiwiteit daarvan verhoog, wat lei tot 'n groter aantal statisties beduidende verskille as wat vasgestel kan word by die gebruik van twee-dimensionele metodes.

Vervolgens stel ons 'n nuwe biologiese visuele analise-metode voor vir die kwalitatiewe analise van kolokalisering. Hierdie metode visualiseer die ruimtelike verdeling van die korrelasie tussen die onderliggende fluoressensie-kanaal-intensiteite met behulp van 'n kleurkaart. Hierdie metode word geëvalueer deur gebruik te maak van sintetiese data sowel as biologiese fluoressensiemikrograwe, wat die verbetering van die visualisering op 'n robuuste manier demonstreer deur slegs waarlik gekolokaliseerde streke aan te dui.

Mitochondriale splitsing, samesmelting en depolarisasie gebeurtenisse is belangrik vir sellulêre funksie en lewensvatbaarheid. Die kwantitatiewe ontleding gekoppel aan die lokalisering van elke gebeurtenis in die drie-dimensionele konteks is egter nog nie gedoen nie. Ons brei die VR-stelsel uit om fluoressensie-gebaseerde tydsverloop-sekwensies van mitochondria te ontleed en stel 'n nuwe metode voor om outomaties die ligging en hoeveelheid van die mitochondriale gebeure te bepaal. Die waargenome liggings van mitochondriale gebeurtenisse kan dan op die fluoressensie-z-stapels aangebring word. Ons pas hierdie metode toe op beide kontrole monsters sowel as selle wat met hydroxychloroquine sulfaat (HCQ) behandel is en demonstreer hoe 'n daaropvolgende kwantitatiewe beskrywing van die splitsing/samesmelting-ewewig sowel as die omvang van depolarisasie bepaal kan word.

Ons kom tot die gevolgtrekking dat virtuele werklikheid 'n aantreklike en kragtige manier bied om fluoressensie-gebaseerde mikroskopie-monsternavigasie, -visualisering en -analise uit te voer. Drie-dimensionele VR-ondersteunde SVB-seleksie laat toe dat monsters met groter noukeurigheid ondersoek en beoordeel kan word en sodoende die potensiaal van fluoressensie-gebaseerde beeldanalise, soos kolokalisering, in biomediese navorsing te benut. Die outomatiese lokalisering en kwantifisering van mitochondriale gebeure kan die navorsing van mitochondriale funksie in gesonde en siek selle ondersteun, waar kwantitatiewe analise van splitsing, samesmelting en depolarisasie van belang is.

Publications

Sections of the work presented in this dissertation have been published as the following research articles in international journals:

- Theart, R.P., Loos, B. and Niesler, T.R., 2017. Virtual reality assisted microscopy data visualization and colocalization analysis. *BMC Bioinformatics*, 18(2), p.64.
- Theart, R.P., Loos, B., Powrie, Y.S. and Niesler, T.R., 2018. Improved region of interest selection and colocalization analysis in three-dimensional fluorescence microscopy samples using virtual reality. *PLOS one*, 13(8), p.e0201965.
- Neethling, A., Engelbrecht, L., Loos, B., Kinnear, C., Theart, R., Abrahams, S., Niesler, T., Mellick, G.D., Williams, M. and Bardien, S., 2019. Wild-type and mutant (G2019S) leucine-rich repeat kinase 2 (LRRK2) associate with subunits of the translocase of outer mitochondrial membrane (TOM) complex. *Experimental Cell Research*, 375(2), pp.72-79.
- Theart, R.P., Loos, B. and Niesler, T.R., 2019. Regression adjusted colocalisation colour mapping (RACC): A novel biological visual analysis method for qualitative colocalisation analysis of 3D fluorescence micrographs. *PLOS one*, 14(11).

Furthermore, the contents of Chapter 6 have been submitted as the following research article.

- Theart, R.P., Kriel, J., Loos, B. and Niesler, T.R., 2019. Mitochondrial event localiser (MEL) to quantitatively describe fission, fusion and depolarisation in the three-dimensional space. (Submitted to PLOS one, January 2020)

Finally, the contents of Chapters 5 and 6 have been used in the following patent applications:

- Theart RP, Loos B, Niesler TR. Regression adjusted colocalisation colour mapping (RACC) - provisional patent filing (VS Ref: P3665ZA00-MS2ST).
- Theart, RP, Kriel, J, Loos, B and Niesler, TR, 2019. Mitochondrial event localiser (MEL) to quantitatively describe fission, fusion and depolarisation in the three-dimensional space - provisional patent filing (VS Ref:P3746ZA00-JC3FA).

Acknowledgements

I would like to express my sincere gratitude to the following people:

- Prof Thomas Niesler and Prof Ben Loos for supervising this project. Their advice, support and encouragement throughout was indispensable for the successful completion of this journey. It was under their guiding hand that I matured as a researcher. Beyond this, their friendship and mentorship played a valuable role in my life. Moreover, their patience in meticulously editing this dissertation.
- Everyone working in the department of physiological sciences at Stellenbosch University who patiently acquired all the samples that I needed for this research. I would especially like to thank Yigael, André and Jurgen.
- Julius Magan, for the spiritual mentorship and fatherly support. Also for the privilege of being part of his ministry for 2 years during the course of this project.
- My family and friends for their support and encouragement.
- My dad, André, for inspiring me to pursue engineering and helping me with conducting the required statistical analysis for this dissertation.
- My mom, Suselna, for her many prayers that carried me through to the end.
- David and Brenda Ramsay, my in-laws, for their love and support. Giving me love, delight and voice.
- My brother, Marnus, for setting the standard of excellence, in whose footsteps I could follow.
- My wife, Jessica, for being my greatest fan and my best friend, always listening to me explaining the many developmental hurdles that I had to face and often giving me the solution without realising it.
- The Lord Jesus Christ for the strength and wisdom to finish this dissertation.

Dedications

I dedicate this dissertation to my beloved wife, Jessica.

Contents

Declaration	i
Abstract	ii
Uittreksel	iv
Publications	vi
Acknowledgements	vii
Dedications	viii
Contents	ix
List of Figures	xi
List of Tables	xviii
Nomenclature	xix
1 Introduction	1
1.1 Background to the project	1
1.2 Region of interest navigation and colocalisation analysis	2
1.3 Mitochondrial fission, fusion and depolarisation analysis	2
1.4 Project objectives	3
1.5 Summary of key contributions	3
1.6 Layout of dissertation	4
2 Base system development and graphics programming background	6
2.1 Base system design objectives and motivation	6
2.2 Graphics programming background concepts	7
2.3 Volume rendering	14
2.4 Virtual reality	26
2.5 Conclusion	28
3 Virtual reality assisted microscopy data visualisation and colocalisation analysis	29
3.1 Introduction	29
3.2 Methods	34
3.3 Results and Discussion	42
3.4 Conclusions	50

4	Improved ROI selection and colocalisation analysis in 3D fluorescence microscopy samples using virtual reality	52
4.1	Introduction	52
4.2	Materials and methods	54
4.3	Results	59
4.4	Discussion	65
4.5	Conclusion	68
4.6	Follow-up work	69
5	Regression adjusted colocalisation colour mapping (RACC) for 3D qualitative colocalisation analysis	70
5.1	Introduction	70
5.2	Methods	73
5.3	Results and discussion	81
5.4	Conclusion	89
6	Mitochondrial event localiser (MEL) to quantitatively describe fission, fusion and depolarisation in 3D	91
6.1	Introduction	91
6.2	Method	93
6.3	Results	100
6.4	Discussion	102
6.5	Conclusion	107
7	Conclusion and future work	109
	Appendices	112
A	Direct volume rendering implementation: Texture-based approach	113
A.1	Overview	113
A.2	Proxy geometry set-up	114
A.3	Conclusion	118
B	Direct volume rendering implementation: Volume ray casting	119
B.1	Overview	119
B.2	Standard ray casting	119
B.3	Pseudo-isosurface ray casting	122
B.4	Conclusion	125
C	MEL synthetic example	126
	Bibliography	130

List of Figures

1.1	The structure of the dissertation. A summary is given of what is contained in each chapter and how these topics related to each other. The green blocks indicate the chapters that were submitted as journal articles. Chapter 2 is the foundation on which the chapters that follow build.	4
2.1	Illustration of the geometric primitives that are available in OpenGL. In all cases, the numbered dots refer to the vertex data (points in 3D space). The primitive names indicate how these vertices should be connected. In the case of a closed polygon, a surface is created.	8
2.2	A simplified representation of the OpenGL graphics pipeline that receives as input a vertex data array and produces the 2D pixels that will be displayed on the screen. The grey blocks are pre-defined and cannot be modified by the programmer. The blue blocks are the shaders that can be modified by the programmer. The fragment shader also receives the texture data which is used to assign a colour to each rasterised fragment. The final block in the pipeline shows the rendered composite image, blended and depth tested with other geometric primitives.	10
2.3	A comparison of the viewing frustum in perspective and orthographic projections (the volume bound between the near and far clip plane). In both cases, the green sphere is outside the frustum and is therefore clipped, while the red and yellow spheres are inside the frustum and are rendered. In the perspective projection, distant objects appear smaller, while in the orthographic projection distant and closer objects have the same size.	14
2.4	Illustration of how a fluorescence microscopy z-stack is stored as a 3D texture in GPU memory, where each block in the 3D texture represents one voxel.	15
2.5	Reconstruction filters for one-dimensional signals. In practice, the box filter (A) or the tent filter (B) are used instead of the ideal <i>sinc</i> function (C) which is not spatially limited.	16
2.6	The light that is emitted at $t = d$ and $t = d_{mid}$ is partially absorbed along the path towards the eye. The accumulated light is blended using a compositing method.	18

2.7	In volume ray casting, a single ray is cast into the volume for each pixel that will be used to make up the final image. For demonstration purposes, we show orthographic projection for the rays which make them perpendicular to the image plane. In perspective projection all the rays originate from the location of the viewpoint. The final colour of this fragment on the bounding box is determined by sampling the volume at regular discrete positions along the ray. These samples are then used to evaluate the volume rendering integral. 1. A ray is cast into the volume for each pixel in the image plane. 2. The volume is resampled at equidistant points. 3. The colour at these sample points is determined using some interpolation scheme or a transfer function. 4. The final colour of the pixel in the image plane is determined by compositing the samples along the ray which assigns the final colour to the fragments on the visible faces of the bounding box.	20
2.8	An illustration of colour bleeding during interpolation. A triangle is drawn with a different colour specified for each vertex. When vertices A and C are specified to be completely transparent, the colour values in the interior of the triangle are interpolated. Even though the vertices are transparent, their colour is still “bleeding” into the interior of the triangle. In the image on the right, colour bleeding was suppressed by pre-multiplying the vertex colours with their opacity before interpolation.	22
2.9	Switching between two texture arrays in the case of object-aligned slices in order to minimise the angles between the viewing direction and slice normals. Between image (C) and (D) the 2D texture arrays were switched.	24
2.10	Object-aligned texture-based volume rendering. A. A parallel projection of samples (yellow) on the object-aligned slices. B. Demonstration of how the distance between adjacent slices depend on the viewing direction which affects the perceived volume density.	24
2.11	Sampling points illustrated for viewer-aligned slicing for orthographic (or parallel) projection (A) and perspective projection (B) respectively. These sampling points are the fragments on the proxy slices which are blended together to create the final image on the image plane. C. Demonstration of how the proxy geometry slices changes as the viewing angle change.	25
2.12	Example of how the IPD relates to the rendering image plane.	27
2.13	Illustration of how distortion affects straight lines. Opposite distortions can be used to rectify a distorted image.	27
3.1	An example graphical user interface (GUI) presented by our system. Here the colocalisation settings GUI is shown. In this case the colocalisation parameters can be adjusted, such as the channels that should be checked for overlap, how the rendering should be performed, what the high and low threshold should be as well as the opacity of the colocalised voxels.	35
3.2	The broader GUI view. The user has GUI panels on both sides that directly face the user when looking in that direction. The colocalisation GUI is to the right, the volume sample in the centre and the results of the colocalisation analysis to the left. The volume sample and results are updated immediately as the parameters in the GUI are varied.	35

3.3	The colocalisation can be rendered in four different ways: a) Overlaying the colocalised voxels as white voxels on the original volume sample. b) Only showing the colocalised voxels in their original fluorescence colours. c) Only showing the colocalised voxels as white. d) Not showing any colocalisation. Note that due to the blue (nuclei) and red (mitochondrial DNA) channels overlapping, the nucleus also appears magenta.	36
3.4	The three implemented ROI selection tools. a) The box selection tool, b) the cylinder selection tool and c) the freehand selection tool. Each tool has been used to select a similar part of the volume sample.	39
3.5	Self-assessment of computer proficiency and prior gamepad experience by the subject. a) Participants' general computer experience, and b) gamepad experience.	41
3.6	The main effects plots for time taken when using gamepad and hand tracking, for a) the gamepad total time, and b) the hand tracking total time for the four categories of self-assessed experience. The dotted line across the plot indicates the overall average time.	43
3.7	Subjective ranking of interfaces. a) The preferred interface to use for colocalisation analysis, and b) the easiest interface to learn to use.	47
3.8	Comparative box plots for objective tests. Box plots for the objective tests that comparing the gamepad (GP) and hand tracking (HT) interfaces.	48
3.9	Comparative interval plots for low and high gamepad experience, showing the mean and 95% confidence intervals. Participants that gave a gamepad experience rating of 1 or 2 were considered to have low experience, while a rating of 3, 4 or 5 were considered as high experience.	50
4.1	WST-1 reductive capacity assay of cells treated with CQ for 6 hours (A) and 24 hours (B), respectively. 100 μ M CQ was sufficient to cause a significant reduction in cell viability after 24 hours of CQ exposure, but not after 6 hours. All data are presented as a percentage of control (Mean \pm SEM). ** ($p < 0.05$ vs control). N = 3 independent experiments.	56
4.2	Comparison of the MIP and the 3D visualisation of some of the main regions of interest analysed for colocalisation. The 3D image shows the selection overlaid on the whole cell, as well as the selection isolated from the remainder of the cell. The depicted cell is from the control sample set ($t = 0$). A: The whole cell. B: The perinuclear region for both a cylindrical and ellipsoidal selection. C: The sampled middle region. D: The sampled membrane region. E: The freehand membrane region (no overlay of ROI shown).	57
4.3	The ROI selection tools employed in the analysis. Combinations of these tools were used to make selections. A: The box tool. B: The cylinder tool. C: The sphere tool. D: The freehand tool.	58
4.4	Showing the whole cell of representative cells for each considered time ($t = 0$, $t = 6$, $t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.	60

4.5	Showing the ellipsoidal perinuclear region of representative cells for each considered time ($t = 0, t = 6, t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.	61
4.6	Showing the freehand membrane region of representative cells for each considered time ($t = 0, t = 6, t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.	62
4.7	Showing some microtubulin strands of representative cells for each considered time ($t = 0, t = 6, t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.	63
4.8	Showing some protein aggregates of representative cell for $t = 6$ with the maximum intensity projection at the top and the 45° 3D view with the ROI isolated from the whole cell at the bottom. Shown from left to right the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels. For the protein aggregates only the results at $t = 6$ were visualised and analysed.	63
5.1	The direct visualisation of the PCC. The relative contribution of each voxel's colocalisation intensity \mathbf{q}_i to the correlation coefficient is determined by the individual terms of the summation in the numerator of Equation 5.3. With x and y axes representing the respective intensities of the two fluorescence channels, we assume a maximum intensity of 255, a mean intensity of 127 and a minimum intensity of 0 for each channel. A) The relative contribution to the PCC is shown for all possible values of \mathbf{q}_i , with green representing a positive and red representing a negative contribution. B) The colour that would be assigned to each voxel with intensity \mathbf{q}_i , using the nMDP colourmap, calculated with Equation 5.1. Note that, at the means, there are sudden changes in the colour assigned by nMDP. C) How nMDP labels each colocalisation intensity \mathbf{q}_i as colocalised or not-colocalised based on its position relative the means. A colocalisation intensity distribution commonly seen for strongly correlated fluorescence channels is shown, with darker shades indicating strong positive correlation and lighter shades indicating weaker positive correlation. The areas marked in red correspond to positively correlated colocalisation that are indicated as negatively correlated when using nMDP.	74
5.2	The projection of colocalised voxels onto the colourmap. A colocalisation intensity \mathbf{q}_i is projected to a 1D subspace defined by \mathbf{p} . A projection onto point \mathbf{p}_0 corresponds to a colourmap value of 0, a projection onto point \mathbf{p}_1 corresponds to a colourmap value of 1, and intermediate projections correspond to a colourmap value varying linearly between these values.	77

- 5.3 The effect of the colour mapping model on possible colocalisation intensities. Using Figure 5.2 as a reference, we consider a specific regression line, shown as a red dashed line, with $\hat{\beta}_1 = 1.5$ and $\hat{\beta}_0 = 0$ for illustration. We visualise the corresponding colourmap value (C_i) calculated for each possible colocalisation intensity $\mathbf{q}_i = (x_i, y_i)$ using the Magma colourmap, where black corresponds to $C_i = 0$ and bright yellow corresponds to $C_i = 1$. For clarity, both thresholds T_{ch1} and T_{ch2} are zero in this example. A: The colourmap values C_i using Equation 5.15. B: The colourmap values C'_i using Equation 5.17 with $\theta = 60^\circ$. C: The final colourmap values C_i using Equation 5.19 with $d_t = 0.2$ and $x_{p_{\max}} = 0.8I_{\max}$ 78
- 5.4 The illustration of the effect of the colourmap thresholds using a synthetically generated scatter plot, where the colourmap brightness indicates the frequency of voxels with a given colocalisation intensity. The regression line is shown in red, with two example intensity pairs, \mathbf{q}_i and \mathbf{q}_j , shown in orange. A penalisation factor θ is applied when projecting both \mathbf{q}_i and \mathbf{q}_j onto the regression line. Only the sample data above the channel thresholds, T_{ch1} and T_{ch2} , are considered when determining the distance threshold d_t , shown in white. Any colocalisation intensity \mathbf{q}_i beyond the maximum point \mathbf{p}_{\max} is assigned a colourmap value of 1. Similarly any \mathbf{q}_i below the minimum point \mathbf{p}_0 is assigned a colourmap value of 0. These thresholds are indicated with two green lines at the respective points. Both d_t and \mathbf{p}_{\max} are determined independently to include 99% of the data, thereby excluding outliers from the visualisation. The final colourmap value C_i varies linearly between 0 to 1 from \mathbf{p}_0 to \mathbf{p}_{\max} 80
- 5.5 The visualisation of synthetic data using both 3D images and maximum intensity projections (MIP). A: The overlapping fluorescence channel intensities. B: All voxels above the colocalisation thresholds and therefore considered colocalised voxels are overlaid in white. C: The result of applying the nMDP. D: The result of applying RACC. E: The difference between RACC and the nMDP. Magenta represents areas which nMDP considers not-colocalised but that RACC considers colocalised, while yellow represents areas that nMDP considers colocalised but RACC does not. F: Scatter plots for the 3D and MIP data sets, showing the frequency of each colocalisation intensity. The regression line calculated by RACC is shown in red. The maximum and distance thresholds for RACC are shown in orange and white, respectively. The red marker represents the per-channel average intensities used by the nMDP calculations and around which the four quadrants are separated (see Figure 5.1). 83
- 5.6 The visualisation of α/β tubulin (red) and acetylated tubulin (green). The structure of the figure is similar to that of Figure 5.5, but in this case biological samples and not synthetic data are considered. 86
- 5.7 The visualisation of autophagosome (green) and lysosome (red) fusion. The structure of the figure is similar to that of Figure 5.5, but in this case biological samples and not synthetic data are considered. 87
- 5.8 The visualisation of autophagosome (green) and tubulin (magenta) interaction. The structure of the figure is similar to that of Figure 5.5, but in this case biological samples and not synthetic data are considered. 88

6.1	Image pre-processing begins by choosing two frames from the input time-lapse sequence, Frame 1 and Frame 2. Normalisation, binarisation, and labelling are then performed on both frames. The labelled frames are each separated into an array of z-stacks, where each stack contains only a single labelled structure. Each of the stacks in the array of binarised stacks, for both Frames 1 and 2, is then Gaussian blurred and Canny filtered. The latter contains only the edges of the labelled structures. The stacks that are used during the automatic image analysis step are labelled in green for Frame 1 and orange for Frame 2. Finally, an array containing the centre of mass for each labelled structure is shown in blue.	94
6.2	Automatic image analysis begins by calculating the overlapping volume of each structure in Frame 1 with each structure in Frame 2. V_{xy} refers to the number of voxels that overlap between structure x in Frame 1 and structure y in Frame 2. To compensate for such coincidental matches, all overlapping volumes that account for less than 1% of the volume of either structure in question are eliminated. Label 0 refers to the background. Using matrix V , the labelled structures that are associated with each other both between Frames 1 and 2, as well as within the same frame, can be calculated. The relationships are encoded as arrays of lists, where each list is related to the labelled structure corresponding to the array index. Therefore, the 8 th list in the array relates to labelled structure number 8. $F2L_8$ refers to label number 8 in Frame 2 that is associated with the structure at the array index. Next, the relative percentage of overlap (e.g. $\%_{F2L_8}$) is calculated for the associated structures between the frames, and is interpreted in conjunction with A_1 and A_2 . The midway points (e.g. \mathbf{p}_4) and distances (e.g. d_4) between two structures within the same frame are calculated, and are interpreted in conjunction with W_1 and W_2 . Using these arrays of lists, the status (e.g. S_{F1L_4}) of fission, fusion, depolarisation or no event can be determined for each mitochondrial structure. Finally, the event can be overlaid on the original z-stack at the location of the midway points.	96
6.3	The back-and-forth structure matching algorithm used to determine mitochondrial structures that are candidates for fission and fusion.	98
6.4	Control sample 1. A. Maximum intensity projection (MIP) of the MEL events overlaid on the entire sample image of every fourth frame in the time-lapse sequence. B. A region of interest (ROI) selection, indicated by the white square in column A. Frame 1 matches the time indicated, Frame 2 shows the subsequent time step, MEL shows the detected mitochondrial events overlaid on Frame 1. C. A selection of the ROI frames in column B, visualised in 3D using volume rendering. D. The five-frame simple moving average of the values calculated by MEL. For the number of events, the area plots indicate the maximum and the minimum that MEL detected for different pre-processing parameters.	102
6.5	Control sample 2. Description as for Figure 6.4.	103
6.6	HCQ treated sample 1. Description as for Figure 6.4.	104
6.7	HCQ treated sample 2. Description as for Figure 6.4.	105
6.8	3D visualisation of MEL event showing the first frame in the time-lapse sequence of the samples in Figures 6.4-6.7. A. Control sample 1. B. Control sample 2. C. HCQ treated sample 1. D. HCQ treated sample 2.	106

A.1	A bounding box is defined around the volume data. Proxy geometry are then generated inside the box. The volume data, stored as a 3D texture, is then resampled over the surface of the proxy geometry slices to create the volume rendering.	114
A.2	An illustration of several arbitrary planes slicing through the bounding box producing proxy geometry polygons consisting of between three and six vertices.	115
A.3	An illustration of a six vertex triangle fan. Vertex 0 is the starting vertex, to which all other vertices connect.	116
A.4	Left: The vertex numbering that is used for the bounding box. There are always three independent paths from the front vertex \mathbf{v}_0 to the back vertex \mathbf{v}_7 , as marked by the solid lines. Right: The plane intersection with the dotted line will always occur between the intersection with two solid lines and must be inserted accordingly.	117
B.1	Ray casting is performed by sampling the volume data set at discrete locations along the ray.	120
B.2	An example of applying pseudo-isosurface ray casting volume rendering to a fluorescence microscopy sample.	123
B.3	An illustration of how the isosurface boundary is determined in pseudo-isosurface ray casting.	123
B.4	An illustration of the Phong shading model. Ambient light refers to the reflected environmental light. Diffuse light refers to the direction impact of light from the light source. Sepecular light refers to the reflective brights spots as they appear on shiny objects.	125
C.1	Synthetic image example process flow.	127
C.2	Synthetic image example calculated matrix and list of arrays.	128
C.3	Example of determining fission and fusion event locations.	129

List of Tables

3.1	Contingency table summary of subjective results. The average rating per question is given on the diagonal. The highest average is highlighted in green, the second highest average is highlighted in orange and the lowest average is highlighted in red. The p-values are shown below the diagonal and the V-values are shown above the diagonal. The p and V-values that are statistically significant are highlighted.	45
3.2	Comparison of objective results. The average times (in seconds) and standard deviation for all the objective tasks for both interfaces are shown, as well as the statistical power values for each task. The p-value from the t-test indicates a statistically significant difference between interfaces when $p < 0.05$. All statistically significant differences are highlighted.	49
4.1	Colocalisation metrics. R_i represents the intensity of voxels in channel 1 (red fluorescent label) and G_i represents the intensity of voxels in channel 2 (green fluorescent label). \bar{R} and \bar{G} represent the mean intensities of the respective channels	56
4.2	Pearson's correlation coefficients calculated for Tau and acetylated α -tubulin in different areas of GT1-7 cells under control conditions and after exposure to chloroquine for 6 and 24 hours respectively. (Data are presented as mean \pm 95% confidence interval). * indicates significant differences relative to control ($p < 0.05$). A shaded block indicates that the ROI selection methods (i.e. 3D vs 2D) are statistically significantly different. In all cases, statistical significance was determined using the two-sample, two-tailed t-test. Protein aggregates were only analysed at the 6 hour time point.	64
4.3	Manders' overlap coefficient with procedure described in Table 4.2.	64
4.4	Manders' correlation coefficient for Tau (M1) with procedure described in Table 4.2.	65
4.5	Manders' correlation coefficient for acetylated α -tubulin (M2) with procedure described in Table 4.2.	65
5.1	Percentage of the total voxels classified differently by RACC and nMDP. . . .	82
6.1	Mitochondrial event percentages for two control samples and two cells treated with HCQ at four different time points corresponding to column A in Figures 6.4-6.7.	101
6.2	The tuneable parameters used by MEL, and the values used for experimentation.	107

Nomenclature

Acronyms and abbreviations

1D	one-dimensional
2D	two-dimensional(-ly)
3D	three-dimensional(-ly)
4D	four-dimensional(-ly)
ALP	autophagy lysosomal pathway
API	application programming interface
ATCC	American type culture collection
ATP	adenosine triphosphate
cFOV	camera field-of-view
Cg	C for graphics
CI	confidence interval
CPU	central processing unit
dFOV	display field-of-view
DMEM	Dulbecco's modified eagles medium
DNA	deoxyribonucleic acid
DVR	direct volume rendering
ETC	electron transport chain
FBS	foetal bovine serum
FITC	fluorescein isothiocyanate
fMRI	functional magnetic resonance imaging
FOC	field-of-view
GFP	green fluorescent protein
GLSL	OpenGL shading language
GP	gamepad
GPU	graphics processing unit
GUI	graphical user interface
HLSL	High-level shading language
HMD	head-mounted display
HT	hand tracking
IMM	inner mitochondrial membrane
IPD	interpupillary distance
KM	keyboard and mouse

LCD	liquid crystal display
LRRK2	leucine-rich repeat kinase 2
LUT	look-up-table
M1	Manders' correlation coefficient (for channel 1)
M2	Manders' correlation coefficient (for channel 2)
MCC	Manders' correlation coefficient / Manders' colocalisation coefficient
MEF	mouse embryonic fibroblast
MEL	mitochondrial event localiser
MIP	maximum intensity projection
MOC	Manders' overlap coefficient
MT	microtubulin
mtDNA	mitochondrial DNA
NDC	normalised device coordinates
NFT	neurofibrillary tangle
nMDP	normalised mean deviation product
OpenGL	Open graphics library
PCC	Pearson's correlation coefficient
PenStrep	penicillin/streptomycin
PINK1	PTEN-induced kinase 1
PONR	point of no return
PSF	point spread function
PTEN	phosphatase and tensin homolog
RACC	regression adjusted colocalisation colour mapping
RGBA	red green blue alpha (opacity)
ROI	region of interest
SDK	software development kit
SMA	simple moving average
SNR	signal-to-noise-ratio
SR-SIM	super resolution structured illumination microscopy
TMRE	tetramethyl rhodamine ethyl
TOM	translocase of outer mitochondrial membrane
VR	virtual reality

Chapter 1

Introduction

1.1 Background to the project

Confocal microscopy is one of the major imaging tools used in molecular life sciences [1]. It utilises lasers and complex illumination settings to excite a fluorescent reagent or probe within a particular focal plane and at a certain defined depth in the z-dimension. This acquisition approach avoids the collection of out of focus light that would originate from above and below the focal plane. Unlike an epifluorescence system, where the whole sample is illuminated at once, confocal microscopy makes use of a pinhole, allowing tight control of illumination in the z-dimension. By exciting the fluorescent reagents at different depths, a z-stack of micrographs (images) is generated. These two-dimensional individual image frames can be reconstructed three-dimensionally using direct volume rendering. In this dissertation, we will use the terms confocal microscopy and fluorescence microscopy interchangeably. The choice of the confocal microscope, linked to super resolution structured illumination (SR-SIM), is motivated by its current high relevance, and the availability of state-of-the-art expertise and equipment at Stellenbosch University.

Current technology in confocal sample visualisation is largely limited to either viewing single micrographs of the z-stack or considering a maximum intensity projection (MIP) of the full z-stack. This is largely due to the limited functionality and feasibility of existing 3D visualisation tools. Up to now, the visualisation of confocal microscopy images has been limited to rendering on two-dimensional displays. The use of three-dimensional displays for both quantitative and qualitative assessment of fluorescence microscopy images has received little attention. This has, consequently, limited the precise selection of regions of interest within the sample, the processing of multiple fluorescence signals to determine their three-dimensional spatial relationship and the analysis of intracellular function in the three-dimensional space.

Recent developments in virtual reality (VR) technology have led to the emergence of headsets that are lightweight while offering high resolution and frame rates, low latency head tracking, and a large field of view [2, 3, 4]. These advances make VR an attractive technology to use in biological analysis. VR allows immersive three-dimensional visualisation, as opposed to a three-dimensional rendering on a two-dimensional display [2]. Because the visualisation is based on a true three-dimensional awareness of the sample, VR offers an unambiguous representation of the data as well as a highly intuitive process of interaction. However, VR guided tools for precision cellular analysis remain exceedingly rare and limited [5, 6].

This project focused on the development of a system for the visualisation and analysis

of three-dimensional fluorescence microscopy data using VR. In addition, new manipulation and analysis methods were developed and specifically implemented for VR, offering unique benefits compared to currently used methods such as unambiguous visualisation and precision selection. The objective of the methods we develop is to improve the current biological understanding of intracellular structures that are implicated in neurodegeneration, such as Alzheimer's disease. To achieve this, we will consider mammalian cells and investigate structures and organelles that are particularly critical for neuronal function. Some of these investigations include the extent of Tau localisation on the tubulin network, the interaction between tubulin and autophagosomes and between autophagosomes and lysosomes, as well as mitochondrial dynamics. The necessary background for these structures will be provided in the relevant chapters.

1.2 Region of interest navigation and colocalisation analysis

In fluorescence microscopy, colocalisation refers to two probes, or colour channels, that codistribute with one another. This can be used to determine whether two proteins or structures associate with one another in the cell. Current tools are limited to calculating colocalisation on either a single micrograph of the z-stack or on a maximum intensity projection (MIP) of the z-stack. Therefore, current colocalisation analysis is limited to two dimensions, not allowing the isolation of defined regions of interest from the surrounding three-dimensional structure before calculating the colocalisation. The region of interest (ROI) is a selected sub-volume of the complete microscopy volume sample that should be considered in a subsequent analysis. For colocalisation metrics to yield meaningful information, it is usually very important to define the region of interest precisely, especially when focusing on particular organelles. Due to the inherent 2D nature of MIP samples, it is not possible to isolate regions of the sample in the 3D space. The need for precise ROI selectivity in all three dimensions is therefore required. Furthermore, to our knowledge virtual reality has never been utilised to perform colocalisation analysis in this manner.

One critical part of the first aim of this project was, therefore, to utilise our new 3D visualisation and analysis tools to perform comprehensive colocalisation analyses using Pearson's correlation coefficient, Manders' overlap coefficient and Manders' correlation coefficients as colocalisation metrics on neuronal cells with detailed VR assisted region of interest selection. We therefore specifically investigated the effect of inhibiting the protein degradation system, i.e. autophagy, on microtubulin and Tau interaction, a molecule component contributing to neuronal degeneration that leads to Alzheimer's disease. This is a question of current research interest in the biomedical field.

1.3 Mitochondrial fission, fusion and depolarisation analysis

Mitochondria can be considered as the "central energetic hub" of the cell, due to their critical role in adenosine triphosphate (ATP) generation through oxidative phosphorylation. ATP availability is evident for cellular function and the onset of cell death, such as apoptosis or necrosis. Recently, it has become increasingly clear that cell death is metabolically regulated and that mitochondria play a more important role in tumour metabolism and

neurodegeneration than previously thought [7]. Therefore, determining mitochondrial function has become a prerequisite for many research groups currently studying the effects of pharmaceutical agents on, for example, cancer metabolism. Even though some 3D analysis of mitochondrial structures exist, the understanding of mitochondrial function remains limited.

To address this, we extend our VR system to localise the mitochondrial events, particularly fission, fusion and depolarisation, in three dimensional space in a time-lapse sequence.

1.4 Project objectives

The overarching objective of this project is to develop a virtual reality based method for the visualisation and analysis of biological samples that can be used to improve the current understanding of cellular function. The specific research context chosen for the application of the developed techniques is neurodegeneration and glioma cells.

In order to improve existing systems and methods currently employed for the analysis and visualisation of confocal image stacks we identified the following research objectives:

1. To design, implement and evaluate a system that can spatially visualise colocalisation in a confocal microscopy sample using real-time interactive virtual reality, as well as calculating several commonly used colocalisation metrics in 3D.
2. To implement different user interfaces appropriate for VR and to determine how intuitive and productive users are in using these interfaces.
3. To design, implement and evaluate various region of interest (ROI) selection tools that allow 3D cell structures to be isolated from the whole sample and visualised by itself in VR. Furthermore, to calculate colocalisation metrics within ROIs selected in this way and to compare these with the commonly used 2D colocalisation analysis.
4. To design, implement and evaluate an improved means of spatially visualising the colocalisation in 3D and to compare this to existing methods by analysing different cellular structures.
5. To design, implement and evaluate an algorithm that automatically determines the location and status of mitochondrial events, specifically fission, fusion and depolarisation, and to visualise the result in 3D.

1.5 Summary of key contributions

The following are the main contributions made by the material contained in this dissertation.

1. A virtual reality enabled system is developed to visualise confocal microscopy samples in 3D by using real-time volume rendering. This system is subsequently applied to the visualisation of colocalisation in 3D, as well as to colocalisation analysis.
2. The productivity, level of control and efficiency, achieved using different input interfaces to the virtual reality system is compared by means of an extensive user trial [8].

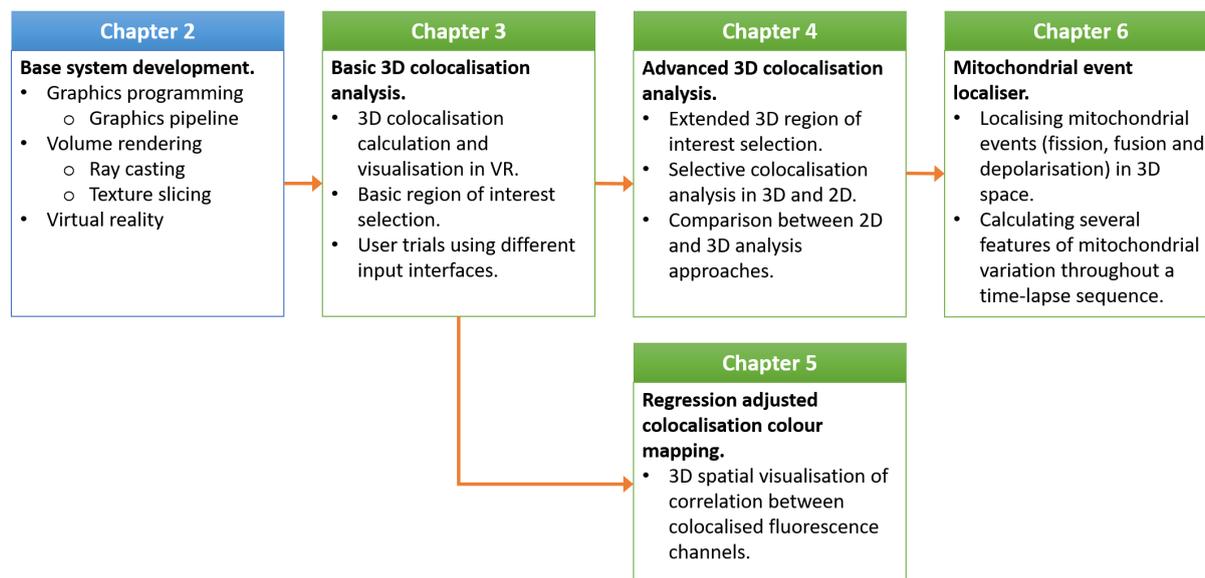


Figure 1.1: The structure of the dissertation. A summary is given of what is contained in each chapter and how these topics related to each other. The green blocks indicate the chapters that were submitted as journal articles. Chapter 2 is the foundation on which the chapters that follow build.

3. The virtual reality system is extended to allow complex and precise 3D region of interest (ROI) selection. By applying the 3D ROI selection tools to neuronal cells, we demonstrate an improvement in the accuracy of colocalisation analysis, compared to corresponding analysis based on projections, leading to new biological insights [9].
4. An approach which allows the visualisation of all colocalised voxels at their 3D location in the sample using a colour map indicating the strength of correlation between two underlying fluorescence channels is proposed. This method is shown to produce results that are easier to interpret and less prone to ambiguous visualisations than current alternatives [10].
5. An algorithm is developed that can determine the 3D location of mitochondrial events, particularly fission, fusion and depolarisation, and visualise these by overlaying the location with the original z-stacks. The algorithm compares two consecutive image frames in a time-lapse sequence and calculates the number of mitochondrial events that occur in each frame, the total number of mitochondrial structures, as well as the average volume of these structures throughout a time-lapse sequence. This method allows the equilibrium between fission and fusion to be indicated numerically, which has not previously been achieved.

1.6 Layout of dissertation

A summary of the layout of the dissertation is shown in Figure 1.1. Since each chapter presents a building block of the overall system and addresses a specific topic, a review of the relevant literature is included in each chapter.

Chapter 2 focuses on the development of the virtual reality (VR) based system that will form the basis of the chapters that follow. This system is developed using volume

rendering and graphics programming. We provide background on these topics to provide a basis for the subsequent chapters.

Chapter 3 extends the VR system, focusing on the visualisation and analysis of colocalisation in 3D. We introduce basic region of interest selection tools and implement different interfaces. We then conduct user trials to evaluate the effectiveness of different input interface methods, and assess the efficiency of performing colocalisation analysis in VR.

Chapter 4 extends the VR system by enabling more complex region of interest selections, as well as the ability to extract a selection and view it in isolation. We apply this process to neuronal cells conducting a detailed colocalisation analysis in 3D and compare the results to a conventional 2D-based analysis.

Chapter 5 extends the colocalisation visualisation further. We propose a novel approach that spatially visualises the correlation between two fluorescence channels. We demonstrate how this improves the understanding of different molecular structures and their association with one another.

Chapter 6 extends the use of the VR system to the analysis of mitochondrial networks. We propose a new algorithm that automatically determines and visualises the 3D location of the mitochondrial fission, fusion and depolarisation events in a time-lapse image sequences. Based on this approach, descriptive features of mitochondria are calculated throughout a time-lapse sequence.

Chapter 7 concludes the dissertation by summarising the main findings and contributions and by providing recommendations for future work.

Chapter 2

Base system development and graphics programming background

This chapter presents the development of a virtual reality (VR) based system that uses direct volume rendering to three-dimensionally (3D) reconstruct and visualise fluorescence microscopy samples. This chapter will lay the foundation on which the remaining chapters will build to answer specific research questions in the life sciences. In the interest of both conciseness and completeness, many details have been relegated to Appendices A and B.

2.1 Base system design objectives and motivation

Fluorescence microscopy samples are generated in the form of z-stacks of two-dimensional images called micrographs. In the following, we developed a visualisation system with the main objective of three-dimensionally reconstructing these z-stacks in order to visualise and interactively manipulate these samples in real-time. Some commercial software, such as ZEISS ZEN, and some open-source tools, such as Fiji, can already achieve this. It is, however, not easy to extend the visualisation and analysis capabilities of these packages. Given our research objectives, it was determined during the early stages of the project that developing our own system, tailored to our specific needs, would provide the necessary flexibility not offered by the available software. One example of this is our desire to use VR to enhance the sample investigation which, at the start of the project, was not possible with any existing systems.

To allow a real-time 3D reconstruction and visualisation from the two-dimensional source slice images, we use several well-established volume rendering methods, each providing particular advantages and drawbacks. Volume rendering is, however, very computationally intensive. To achieve hardware-accelerated rendering, we use a *graphics processing unit* (GPU). This allows certain tasks that can be executed in parallel to be performed much faster than executing them on a general-purpose *central processing unit* CPU.

2.1.1 GPUs and CPUs

GPUs are designed and optimised to perform the same operation on a large batch of data in parallel at very high speeds, relative to a CPU, which focuses on many different tasks and operations at any given time. Architecturally, a CPU consists of a small number of cores that can handle a small number of software threads at a time, and are optimised for sequential serial processing. A GPU, in contrast, has a massively parallel architec-

ture consisting of thousands of smaller cores that are designed to handle thousands of threads simultaneously [11]. With these multi-threading abilities, some software can be accelerated to run more than 100 times faster on a GPU than on a CPU [12]. With the use of GPU-accelerated computing, the performance of applications can be improved by offloading computationally intensive but repetitive portions of the code to the GPU, while the remainder of the sequential code runs on the CPU.

Since volume rendering requires large data sets to be processed in a repetitive way, the highly parallel nature of a GPU enables this data to be rendered in real-time which allows samples to be investigated interactively.

2.1.2 Graphics APIs and Unity

A graphics *application program interface* (API), such as OpenGL or Direct3D, allows the programmer to interact with a GPU at a fairly high level using a *shading language*, without having to be familiar with the assembly instructions for the specific GPU [13]. Different shading languages are available for use in the graphics APIs, the most common being OpenGL's GLSL and Direct3D's HLSL. These languages allow the programmer to control how the GPU renders a 3D scene to a 2D screen. This process is commonly referred to as the *graphics pipeline*, and will be described in Section 2.2.4.

We have implemented our visualisation system in the Unity real-time development platform [14]. One of the main shading languages that Unity uses is Cg (C for graphics), which was developed by Nvidia in close collaboration with Microsoft. Even though Cg is implemented for several graphics APIs, including OpenGL and Direct3D, we will use the OpenGL API with Cg code snippets when describing some general graphics programming concepts which are required to explain how our visualisation system renders the fluorescence microscopy samples.

2.2 Graphics programming background concepts

This section presents a brief overview of some of the main concepts and terminology which are used in graphics programming. We will use the OpenGL graphics API and Cg as the shading language.

2.2.1 OpenGL

The *Open Graphics Library* (OpenGL) is a graphics API for accessing features in graphics hardware that are typically used for 3D graphics rendering. The OpenGL specification includes a description of high-level shading languages, such as GLSL and Cg, and is managed and maintained by the Khronos Group, a non-profit technology consortium.

OpenGL is designed to be a hardware-independent interface that can be implemented on a large variety of different graphics systems, or entirely in software. It is also independent of the operating system or windowing system. As a result, OpenGL does not include functions to perform windowing tasks, process user input or read images, since these are specific to the operating system that is used. For these functions, we used Unity's built-in API. Similarly, OpenGL does not provide functionality for describing three-dimensional objects. Instead, three-dimensional objects are constructed from a small set of *geometric primitives* such as points, lines and triangles [15].

2.2.2 Geometric primitives

A geometric primitive is a basic 3D drawing element. It is a one- or two-dimensional entity, such as a line or a triangle, that is positioned in three-dimensional space and assembled in order to create a 3D object [13]. In OpenGL there are several standard geometric primitives which can be specified using *vertex data* (see Figure 2.1). Vertex data, or *vertices* for short, refer to points in 3D space. The geometric primitives that should be used to connect the vertices are specified when the rendering call is made. The primitive type, together with the vertex data, is passed to the vertex shader (Section 2.2.4).

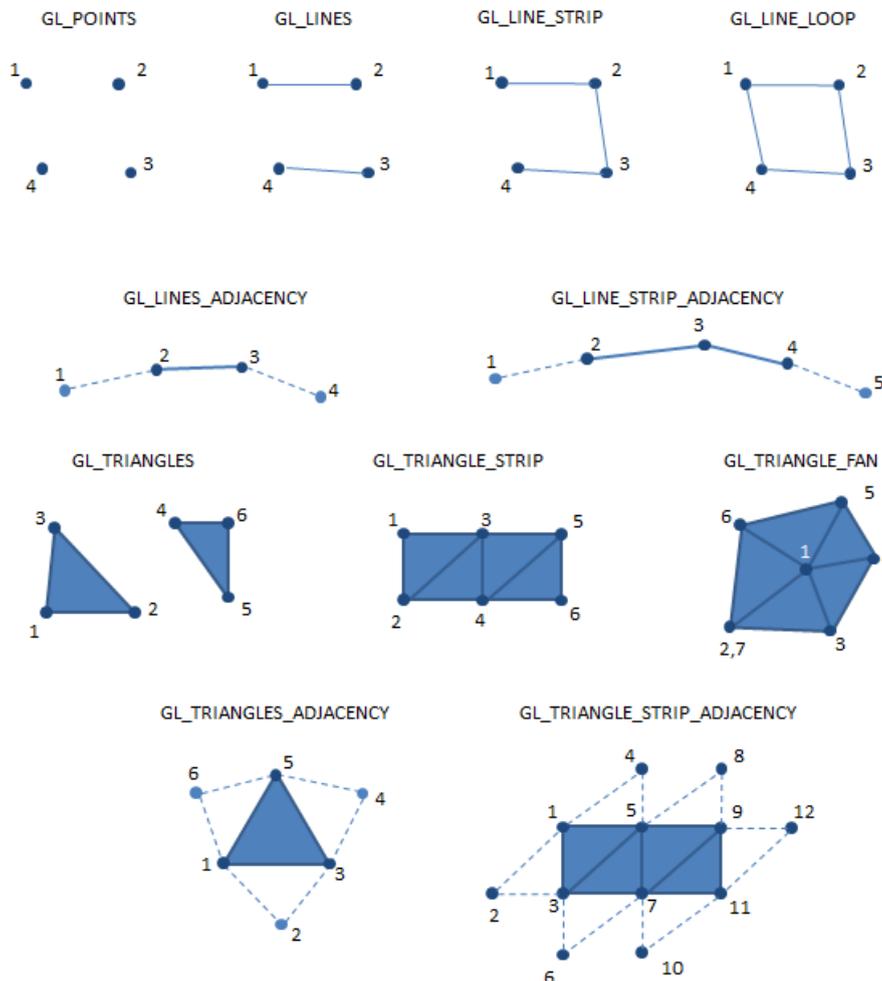


Figure 2.1: Illustration of the geometric primitives that are available in OpenGL. In all cases, the numbered dots refer to the vertex data (points in 3D space). The primitive names indicate how these vertices should be connected. In the case of a closed polygon, a surface is created. Adapted from Neider *et al.* [16].

Using these geometric primitives, it is possible to generate polygons in 3D space. The surface of these polygons can then be overlaid with a colour or an image using *texture mapping*.

2.2.3 Texture mapping

When adding colour to 3D models, it is possible to assign a different colour to every vertex. The colour of the 3D surface is obtained by interpolating between the vertex colours. This is, however, inefficient if image detail needs to be overlaid onto 3D models. For these cases, we make use of a process termed *texture mapping*.

A *texture map*, or simply a *texture*, refers to a 1D, 2D or 3D array of data that is stored in GPU memory. It should be noted that a 3D texture is not a volumetric primitive. Instead, it is a 3D look-up-table (LUT) which should be used in association with the vertex data defining geometric primitives. There is no restriction on the data that can be stored in a texture [15]. However, in our case the textures will be the 2D images in the acquired z-stack, containing the fluorescence channel intensities.

Textures are composed of *texels*, a contraction of “texture elements”, and is the texture-based counterpart of pixels (picture elements) in images. In the case of an image (2D texture), each texel contains colour values. There is, however, utility in considering a texture simply as a LUT. This LUT can then be accessed from within the shaders (Section 2.2.4), especially fragment shaders, where the resulting colour of each pixel in the final rendered image can be determined by looking up the corresponding value in the texture. This value is looked-up using *texture coordinates*, which are values between 0.0 and 1.0 for each dimension, where (0.0, 0.0, 0.0) is associated with the front-bottom-left of a 3D texture, and (1.0, 1.0, 1.0) refers to the back-top-right of a 3D texture.

Whenever a one-dimensional texture is used in a shader and a texture coordinate is requested that lies in-between texels, a weighted average (*linear interpolation*) of the texels surrounding the requested texture coordinate is computed [13]. When a 2D texture is used, *bilinear interpolation*, an extension of linear interpolation to two independent variables (x and y), is used. In the case of 3D textures, this is further extended to three dimensions, x , y and z , and is referred to as *trilinear interpolation*.

2.2.4 The graphics pipeline

The *graphics pipeline*, sometimes also called the *rendering pipeline*, refers to the sequence of steps that leads to a 2D rasterised (array of pixels) representation of a 3D scene. The pipeline is followed for every frame to be rendered. The OpenGL graphics pipeline can be divided into two main parts: first the 3D coordinates are projected onto a 2D coordinate plane and second the 2D plane is transformed into the final array of coloured image pixels that can be displayed on a screen.

The graphics pipeline is divided into several stages where each stage takes the output of the previous stages as input. Since each stage has one specific function, they can be executed in parallel. This inherently parallel nature of the rendering process is one reason why GPUs can accelerate rendering by utilising thousands of small processing cores. Many small programs, called shaders, run simultaneously to accomplish the rendering operation [17].

Shaders¹ are small programs, written by the programmer in a C-like shading language such as GLSL or Cg, which are responsible for the various stages in the graphics pipeline, shown in Figure 2.2. In the figure, the blue blocks represent the stages that the programmer can configure by supplying the code that implements vertex and fragment shaders. The OpenGL graphics pipeline allows for several more shaders, such as the tessellation

¹The name *shader* is a historical name which is easily misinterpreted, since modern shaders do much more than just shading the rendered object.

and geometry shaders. However, these shaders have default (pass-through) versions that will be executed when not supplied by the programmer. Since these shaders were not implemented in our visualisation system, they will not be described further. Currently, OpenGL does not allow the programmer to supply shaders for the greyed-out stages in the pipeline [17]. All the individual parts of the graphics pipeline will now be discussed briefly, with reference to Figure 2.2 as well as Listing 2.1.

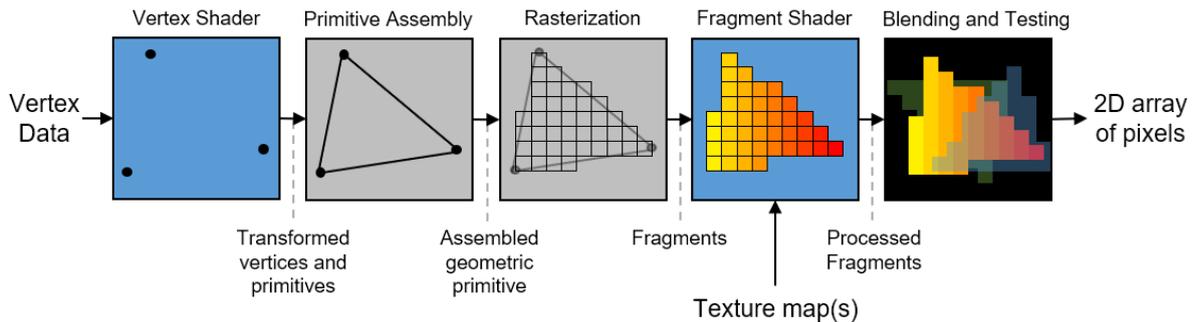


Figure 2.2: A simplified representation of the OpenGL graphics pipeline that receives as input a vertex data array and produces the 2D pixels that will be displayed on the screen. The grey blocks are pre-defined and cannot be modified by the programmer. The blue blocks are the shaders that can be modified by the programmer. The fragment shader also receives the texture data which is used to assign a colour to each rasterised fragment. The final block in the pipeline shows the rendered composite image, blended and depth tested with other geometric primitives. Adapted from De Vries [17]

In Listing 2.1. we define a 2D texture in line 7 and create the associated `sampler2D` variable in line 37. This is simply a 2D array, usually an image, that is stored in GPU memory. However, for this example, we simply made it the colour red. In lines 23-35 we define the input and output data structures for the vertex shader. The output of the vertex shader is then passed to the fragment shader. The vertex and fragment shaders are defined in lines 40-54.

```

1 Shader "SampleShader"
2 {
3     Properties
4     {
5         // The 2D texture that is stored in GPU memory
6         // the data is loaded from the main program
7         _MainTex ("Texture", 2D) = "red" {}
8     }
9     SubShader
10    {
11        Tags { "RenderType"="Opaque" }
12        LOD 100
13
14        Pass
15        {
16            CGPROGRAM
17            #pragma vertex vert
18            #pragma fragment frag
19
20            #include "UnityCG.cginc"

```

```
21
22     // vertex data received from the main program
23     struct dataFromApp
24     {
25         float4 vertex : POSITION;
26         float2 uv : TEXCOORD0;
27     };
28
29     // the transformed vertex and texture coordinates
30     // sent from the vertex shader to the fragment shader
31     struct vertexToFragment
32     {
33         float2 uv : TEXCOORD0;
34         float4 vertex : SV_POSITION;
35     };
36
37     sampler2D _MainTex; // the 2D texture
38
39     // vertex shader
40     vertexToFragment vert (dataFromApp v)
41     {
42         vertexToFragment o;
43         o.vertex = UnityObjectToClipPos(v.vertex);
44         o.uv = TRANSFORM_TEX(v.uv, _MainTex);
45         return o;
46     }
47
48     // fragment shader
49     fixed4 frag (vertexToFragment i) : SV_Target
50     {
51         // sample the colour from the texture at the UV coordinates
52         fixed4 colour = tex2D(_MainTex, i.uv);
53         return colour;
54     }
55     ENDCG
56 }
57 }
58 }
```

Listing 2.1: Basic example shader written in the Cg shading language.

Vertex data

A vertex is a point in 3D space at which two or more straight line segments meet. OpenGL requires all vertex data to be stored in so-called *buffer objects*, which are sections of memory on the GPU that are managed by OpenGL. The vertex data usually consists of a list of 3D coordinates that are grouped to form one of a few OpenGL primitives, such as a triangle, a point or a line (see Section 2.2.2) [17]. Other than the positional vertex data, it is common to include other values in the input buffer objects which are used to determine the rendered pixel's final colour, such as texture coordinates or vertex colours. The buffer is then passed as input to the vertex shader with an OpenGL function call, at which point the geometric primitive is also specified. The input data is grouped in a data structure such as the one shown in lines 23-27 of Listing 2.1.

Vertex shader

As shown in Figure 2.2, the first stage in the graphics pipeline is the vertex shader. This receives a vertex buffer stream as input, one vertex at a time. The main purpose of the vertex shader is to transform the received 3D coordinates into a new 3D coordinate system based on what should be visible in the final image [17]. This is usually performed as a series of matrix multiplications. In Unity this transformation is accomplished by the `UnityObjectToClipPos()` function as shown in line 43 of Listing 2.1.

A secondary purpose of the vertex shader is to update the texture coordinates (see Section 2.2.3) that are used later by the fragment shader. During the rasterisation stage, these coordinates are interpolated over the surface of the primitive, to determine each fragment's associated texture coordinate. The transformation of the texture coordinates is performed in line 44 of Listing 2.1 using the `TRANSFORM_TEX` macro which ensures that the texture is correctly scaled and offset for the current vertex position. This produces a 2D UV coordinate, where the U and the V denote the axes of the 2D texture used by the fragment shader. The x , y and z coordinates denote the axes of 3D world space in which the vertex data is specified.

Primitive assembly

The primitive assembly stage receives the processed vertices from the vertex shader, as well as the vertex connectivity information specified in the main program. These vertices are then connected to create a geometric primitive, such as a triangle (see Section 2.2.2). This stage is always carried out internally in OpenGL [16].

Clipping

This stage in the OpenGL pipeline (not shown in Figure 2.2) checks if the assembled primitives are inside the so-called *clipping volume*. This is the region inside the *view frustum* (refer to Section 2.2.5) that is to be rendered. Primitives that are fully inside the clipping volume are not altered, while those that are outside are discarded. Primitives whose edges intersect the clipping volume are clipped, by creating new primitives whose edges lie on the border of the clipping volume [16].

Rasterisation

After clipping, the rasterisation stage maps the resulting primitive, a 3D vector object, to the corresponding pixels of the final image. For each of these pixels, a *fragment* is generated which contains the raster position, depth value as well as interpolated values such as colour, opacity and texture coordinates associated with the particular fragment. These fragments are then sent to the fragment shader, where they are individually assigned a final colour and opacity [17].

Fragment shader

The main purpose of the fragment shader is to calculate the final colour of each pixel within the rendering window, that will be displayed on the screen. The fragment shader is executed for each of the fragments generated during the rasterisation stage and receives its transformed x , y and z coordinate, where the latter is used when performing depth testing (see the next section). The shader also receives the interpolated UV coordinate that is used to sample from the 2D texture map (see Section 2.2.3), as is shown in line 52 of Listing 2.1. The only requirement for a fragment shader is that it must return a four-dimensional vector, containing the fragment's colour as a Red Green Blue Alpha

(RGBA) value, with alpha representing the opacity. In Listing 2.1 this output vector is called `colour`.

Blending and depth testing

In the final stage of the OpenGL pipeline, the visibility of a fragment is determined using *depth testing* (also referred to as *z-buffering*). This determines whether there are any fragments in front of the fragment under consideration. If a fragment passes the depth test, it can be directly written to the *framebuffer*, which stores the final colour of the rendered pixel. However, if *alpha blending* is enabled, this stage also considers the alpha value. Alpha values define the opacity of a fragment and are usually the fourth value in a colour vector. The colour of the fragments behind the current fragment are blended and the resulting colour is written to the framebuffer [15]. When the alpha value is 0.0, the fragment is completely transparent, and when it is 1.0 the fragment is completely opaque. Semi-transparency is achieved with values between 0.0 and 1.0.

Framebuffer

The framebuffer stores the fragments that can be directly displayed on the screen by the GPU hardware. It is, in fact, a combination of several buffers, including the colour buffer, which stores colour values, and the depth buffer that stores depth information for each pixel [17]. The framebuffer is the destination of the rendering process. The fragment shader writes the fragment information to the framebuffer, from where the window manager displays it. Note that blending and depth testing occurs when a new fragment is written to a framebuffer location that already has a value due to it being the destination of a previous fragment.

2.2.5 Coordinate systems and projections

The coordinate system of the 3D scene is transformed in the vertex shader by a series of matrix multiplications, resulting in a scene that uses the coordinate system required by OpenGL, known as *normalised device coordinates* (NDC). One intermediate step when transforming the scene into NDC is to transform the scene into the *view space*, which is often also referred to as the *camera*, *camera space* or *eye space*. This view represents what is seen from the camera's, or the user's, point of view.

The coordinate system that follows the transformation into the view space is called the *clip space*, which determines which parts of the scene should be visible and which parts are outside of the visible range and should, therefore, be *clipped* or discarded. During this transformation the range of the visible coordinates is termed the *frustum* and every coordinate inside the frustum will be rendered to the screen [18]. The shape of the frustum determines the type of projection that will be used. The two types of projections that are typically used are the *perspective projection* and an *orthographic projection* (both are illustrated in Figure 2.3). In both cases, a near and a far clip plane are defined. The near clip plane is usually associated with the *image plane* which is rendered to the screen and the far plane refers to the distance that a 3D object can be away from the viewpoint and still be visible. With an orthographic projection, a width and a height must also be specified. Together with the near and far clip planes, these are used to define a cube-like frustum in which distant objects appear the same size as closer objects. To render objects the way they appear in the real world, where distant objects appear smaller than closer objects, a perspective projection must be used. Instead of specifying a width and height,

a perspective projection requires a *field-of-view* (FOV) to be specified [17]. The FOV is an angle describing the extent of the observable world that is seen at any given time. Usually, only the vertical FOV is specified. The horizontal FOV is then derived from the rendering window's aspect ratio.

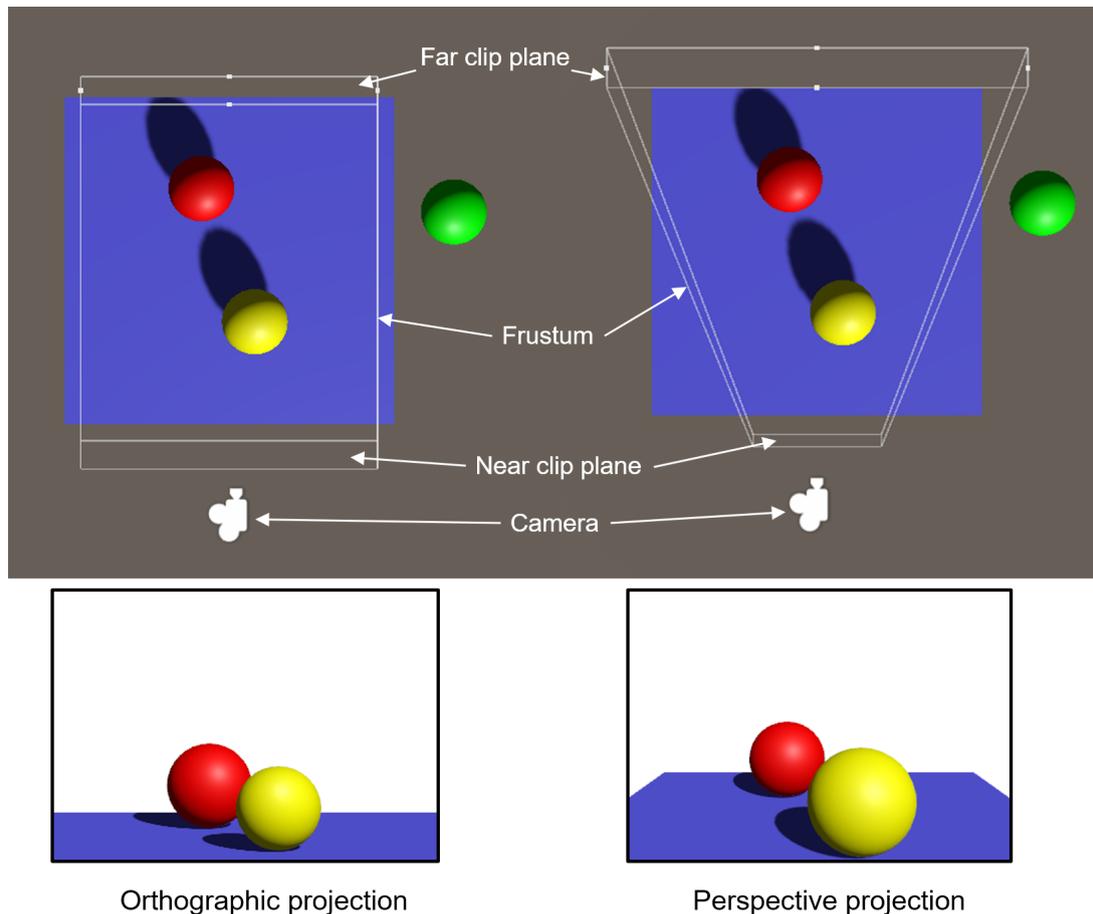


Figure 2.3: A comparison of the viewing frustum in perspective and orthographic projections (the volume bound between the near and far clip plane). In both cases, the green sphere is outside the frustum and is therefore clipped, while the red and yellow spheres are inside the frustum and are rendered. In the perspective projection, distant objects appear smaller, while in the orthographic projection distant and closer objects have the same size. Adapted from Rougier [19]

2.3 Volume rendering

Confocal microscopy produces a three-dimensional scan of a biological sample by capturing multiple two-dimensional micrographs to produce a so-called *z-stack* of slice images. This *z-stack* can be three-dimensionally reconstructed by means of *volume rendering*. The term volume rendering describes a wide range of techniques that can be used to project a 3D scalar data set onto a 2D plane from any angle. The projections is performed in such a way that it creates the illusion of a solid 3D structure. Other volume rendering methods extract an isosurface (surface of constant value) which could take the form of a

polygonal mesh [20, 21, 22]. Modern volume rendering techniques take advantage of the superior performance of graphics accelerators to achieve real-time data visualisation [21].

A typical three-dimensional volumetric data set consists of a series of 2D equispaced slice images where each image contains the same number of pixels in a regular grid pattern. Such a 3D data set represents a discretely sampled regular volumetric grid of volume elements, known as voxels (illustrated in Figure 2.4). Each voxel is associated with one or more constant data values representing attributes such as opacity and colour. Unfortunately, the term *voxel* has two slightly different interpretations in the literature. One interpretation is that voxels are small cubes that each fill a volumetric region. The other interpretation is that voxels are points in 3D space, which are tied to an interpolation scheme that fills the space between the points. The latter is the interpretation followed in this discussion of volume rendering, since it allows for a more flexible reconstruction of in-between data values, which is a requirement for high-quality volume rendering [20]. One major advantage of using modern graphics hardware is its ability to perform very fast trilinear interpolation between voxel data.

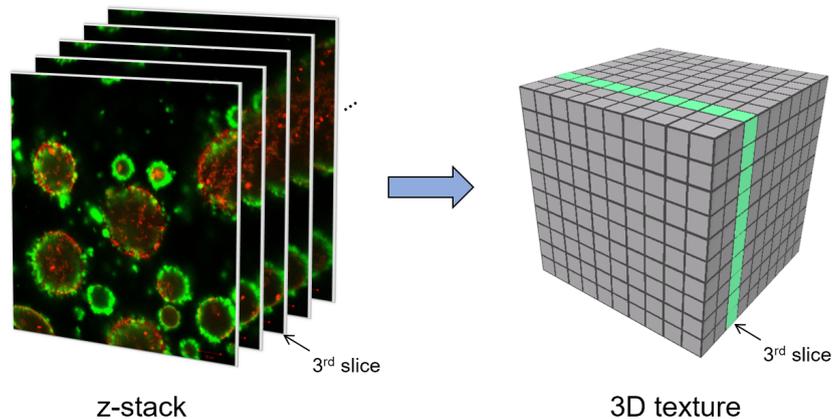


Figure 2.4: Illustration of how a fluorescence microscopy z-stack is stored as a 3D texture in GPU memory, where each block in the 3D texture represents one voxel.

2.3.1 Reconstruction from voxel data

Even though it may be more intuitive to imagine voxels as tiny cubes for volume rendering, it is more accurate to think of each voxel as a sample, such as the intensity of a fluorescence channel, that is obtained at a single infinitesimally small point in a continuous three-dimensional signal space as represented by:

$$f(\vec{x}) \in \mathbb{R} \quad \text{with} \quad \vec{x} \in \mathbb{R}^3 \quad (2.1)$$

where \vec{x} represents the sampled voxel.

Provided that the continuous signal f is spatially band-limited and has a cut-off-frequency of F_{max} , according to the sampling theorem, exact reconstruction is possible if the signal is sampled at a rate more than two times the cut-off-frequency, or $F_s > 2F_{max}$. This border sampling frequency of $2F_{max}$ is also known as the Nyquist rate [23]. In practice, there are however two major problems which prevent perfect reconstruction of the sampled volume data.

1. For perfect reconstruction, according to the sampling theorem, a convolution of the sample points with a *sinc* function (Figure 2.5c) is required in the spatial domain. In the one-dimensional case the *normalised sinc function* is defined as

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}. \quad (2.2)$$

The three-dimensional case is obtained by using the tensor-product. Note that the function has an infinite extent. Therefore, to perfectly reconstruct the original signal at an arbitrary position, all the sampling points must be considered. In practice, this is not computationally feasible with limited hardware and time [20].

2. Generally the volumetric data does not represent a band-limited signal. Sharp boundaries between different materials are spatial step functions which are not band-limited. The sampling and reconstruction of a signal that is not band-limited will result in aliasing artefacts [20].

Therefore, to reconstruct a continuous signal in practice from an array of voxels the ideal 3D *sinc* filter is usually replaced by either a box filter (Figure 2.5a) or a tent filter (Figure 2.5b). When using a box filter, nearest-neighbour interpolation is applied which results in a blocky appearance of the output, which is usually undesirable. Most modern graphics cards automatically apply trilinear interpolation, which is implemented as a convolution with a 3D tent filter. This represents a trade-off between computational cost and output quality [20].

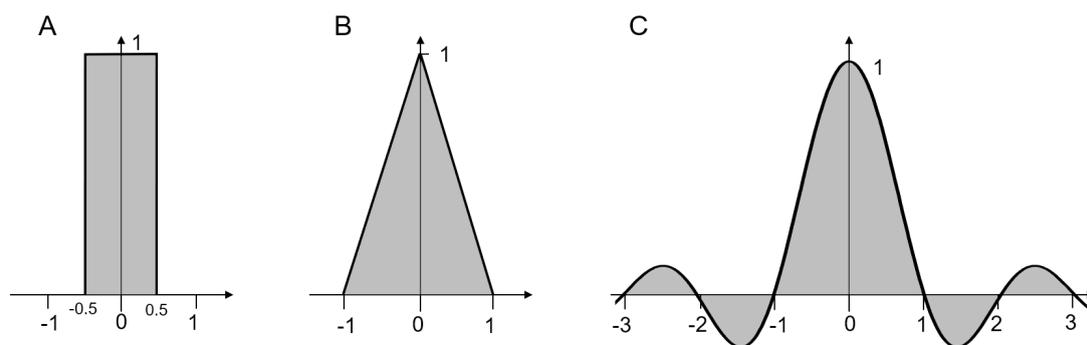


Figure 2.5: Reconstruction filters for one-dimensional signals. In practice, the box filter (A) or the tent filter (B) are used instead of the ideal *sinc* function (C) which is not spatially limited. Adapted from Engel *et al.* [20].

2.3.2 Direct volume rendering

Direct volume rendering (DVR) is a set of techniques that generate 2D projected images from a 3D volumetric data set, in our case the z-stacks obtained by fluorescence microscopy. The z-stack data is loaded as a 3D texture map into the memory of the GPU where it is rendered, without explicitly extracting surface geometry from the data [24]. This stands in contrast to indirect volume rendering, where the surface geometry mesh is extracted before rendering takes place. Indirect volume rendering is less applicable to the rendering of fluorescence microscopy images, because it usually produces an isosurface, thereby removing all internal structural information. Fluorescence microscopy data analysis is often concerned with how different molecular structures overlap, which can best be

represented using direct volume rendering methods where all internal structural overlap can be more clearly discerned.

DVR techniques make use of an *optical model* to map generated or measured voxel data to *optical properties*, such as colour and opacity [25]. During rendering, these optical properties are composited (blended together) in a predefined way to produce a visualisation of the data. Some of these compositing methods will be discussed in this section.

The optical model describes how the particles in the volume react with light. There are several variations of the optical model. One of the most common models assumes that the volume consists of particles that both emit and absorb light. In other words, every voxel is modelled to emit coloured light but also has an opaqueness causing light from other voxels to be absorbed. The optical parameters are either specified directly by the volume data set or are calculated by applying one or more *transfer functions* to the data set [21]. It is important to note that these optical models and transfer functions are usually not intended to create a photo-realistic rendering of the subject, they are simply a means of visualisation.

A transfer function is used to either emphasise or to clarify certain features of interest in the volume. A volume data set usually contains abstract scalar data and generally there is no natural way to obtain the emission and absorption parameters. Therefore, with the use of a transfer function, the user specifies what the different structures in the volume should look like by specifying an arbitrary mapping from the 3D texture map values to the rendered fragment colours [20]. Transfer functions are used to apply this mapping and are typically implemented as a *texture lookup table* or as a simple function within the fragment shader.

In the case of fluorescence microscopy, different fluorescent probes are used that bind to different molecular structures. These probes have a natural fluorescence colour and these same colours will be used in the transfer functions to visualise fluorescence intensities. Furthermore, in Chapters 3, 4 and 5 a different transfer function will be used to visualise regions where the two fluorescent probes overlap in the same voxel, a phenomenon that is known as colocalisation. Lastly, in Chapter 5 we will develop a novel transfer function that highlights regions of strong positive correlation between the fluorescence channels.

2.3.2.1 The volume rendering integral

When rendering volumetric data directly, the volume is modelled as being composed of a semi-transparent material that simultaneously emits and absorbs light. The intensity of the light as it passes through the volume is computed using the *volume rendering integral* [25].

The volume rendering algorithm that evaluates this integral most directly is known as volume ray casting (see Section 2.3.2.2). In the simplest case, volume ray casting passes a single ray through the volume, originating from the viewpoint, and integrates the optical properties of all voxels along the ray's path [20]. The ray's direction is determined by the compositing method that is used which could either be in front-to-back or back-to-front order. These two directions produce the same visualisation result, however the front-to-back order offers some implementation specific optimisations which will be discussed later. It is useful to keep this volume rendering method in mind when deriving the volume rendering integral.

We will present a simple volume rendering equation based on an emission-absorption optical model, as was derived by Engel *et al.* [26]. We define the ray that is cast into the volume as $\vec{x}(t)$, where t is the distance from the viewpoint, or the eye, which corresponds

to the image plane. The scalar value, representing the volume's optical properties, that corresponds to a position along the ray, can then be defined as $s(\vec{x}(t))$. Now we define the *absorption coefficient*, which accounts for the absorption of light along the ray, as $\kappa(s)$, and the *emission coefficient*, which accounts for the radiant energy that is actively emitted along the ray, as $c(s)$. For simplicity we will define the emission and absorption coefficients, c and κ as functions of the distance t to the eye:

$$c(t) := c(s(\vec{x}(t))) \quad \text{and} \quad \kappa(t) := \kappa(s(\vec{x}(t))) \quad (2.3)$$

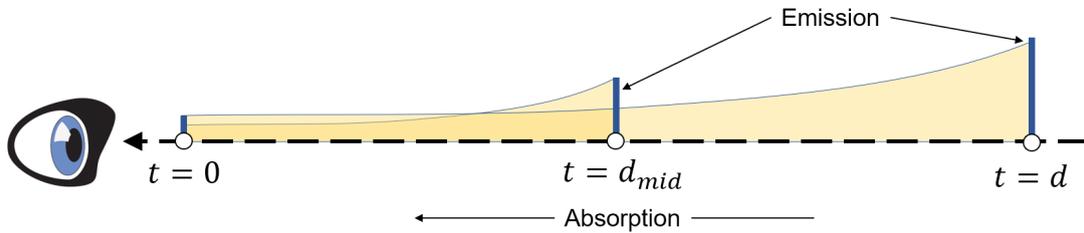


Figure 2.6: The light that is emitted at $t = d$ and $t = d_{mid}$ is partially absorbed along the path towards the eye. The accumulated light is blended using a compositing method. Adapted from Engel *et al.* [26].

The idea of emission and absorption is illustrated in Figure 2.6. Light is emitted at position $t = d$ along the ray and is continuously absorbed along the length of the ray until it reaches the eye. Therefore, only a certain portion c' of the light emitted at $t = d$ will reach the eye. By adding more points of emission along the ray, such as the point at $t = d_{mid}$, the amount of light that reaches the eye will accumulate using some blending algorithm. If constant absorption ($\kappa = \text{constant}$) along the ray is assumed, c' will be:

$$c' = c \cdot e^{-\kappa d} \quad (2.4)$$

In the case where κ is not constant along the ray but is also dependent on the position t , the amount of light c' that reaches the eye can be calculated by integrating the absorption coefficient along the ray:

$$c' = c \cdot e^{-\int_0^d \kappa(\hat{t}) d\hat{t}} \quad (2.5)$$

For the general case, the absorption coefficient integral can be defined as:

$$\tau(d_1, d_2) = \int_{d_1}^{d_2} \kappa(\hat{t}) d\hat{t} \quad (2.6)$$

In Equation 2.6, τ is known as the *optical depth*, and $\tau(0, t)$ can be substituted into Equation 2.5.

Up to this point, we have assumed that light was only emitted at a single point along the ray. To calculate the total light or radiation energy C that reaches the eye along a particular ray, emissions from all possible positions t must be considered:

$$C = \int_0^{\infty} c(t) \cdot e^{-\tau(0,t)} dt \quad (2.7)$$

Equation 2.7 is known as the *volume rendering integral*, where C can be thought of as the intensity of the light that the eye sees. This intensity can be remapped to a colour

and opacity for final rendering by using a transfer function. In the case of fluorescence microscopy, the volume rendering integral is calculated for each colour channel that is present, usually red, green, blue and magenta, after which the four resulting values of C are blended together.

In practice, the volume data set is discretised from a continuous source (see Section 2.3.1), and hence the integral in Equation 2.7 is usually approximated by an appropriate Riemann sum. This is then evaluated numerically in either back-to-front or front-to-back order with alpha blending composition, as illustrated by the volume ray casting method, which we consider next.

2.3.2.2 Volume ray casting

Volume ray casting is an *image-order* volume rendering technique. This means that the rendering algorithm iterates over the pixels of the final rendered image to produce the visualisation of the data. It is important not to mistake volume ray casting, which processes volume data, with ray tracing, which processes surface data. In volume ray casting, the computation does not stop at the surface. Instead, the rays pass through the volume to determine the resulting colour and opacity of rendered pixels. Furthermore, reflections do not lead to secondary rays, as they do in ray tracing [27].

Volume ray casting, from here on referred to simply as ray casting, can be derived directly from the volume rendering integral (Equation 2.7). When a 3D data set is visualised, it is common to define a rectangular prism, known as the bounding box, that defines the limits of the volume. When this bounding box is rasterised (Section 2.2.4), fragments are generated over its visible faces. For each of these fragments, a single ray is projected into the volume (either in back-to-front or front-to-back order). The volume data is then resampled at regular discrete positions along the ray, usually using trilinear interpolation as the reconstruction filter. These samples are blended to determine the final displayed image. The volume is, therefore, projected onto the surface of the bounding box in such a way as to create the illusion of the 3D reconstructed volume occupying the space inside the bounding box. This process is illustrated in Figure 2.7.

The natural direction of transversal through the volume is front-to-back since rays conceptually originate at the camera [20]. The camera represents the location from which the 3D scene is viewed (the eye). If the volume is sampled in front-to-back order, it allows for *early ray termination*, where the calculation stops as soon as the opacity along the partial path is greater than a predetermined threshold. This eliminates computation that does not contribute substantially to the rendered image. Computational efficiency can also be improved by a method called *empty space skipping*. In this method, transparent samples along the ray, as determined by a transfer function, are skipped [26].

Ray casting produces a very high-quality rendering result. It was used in the past as a CPU volume rendering method since early GPUs did not support the functionality required for its implementation. With the introduction of fragment shaders, however, it is now possible to perform GPU-based volume ray casting [20, 22].

Discrete approximation of the volume rendering integral

In order to implement the volume rendering integral (Equation 2.7) in a fragment shader, such as is required for ray casting, it must first be approximated using a Riemann sum. We will present that process in this section as was derived by Hadwiger *et al.* [29].

The optical depth τ (Equation 2.6), is the cumulative absorption of radiation energy along the ray up to a certain position $\vec{x}(t)$, can be approximated by a Riemann sum by

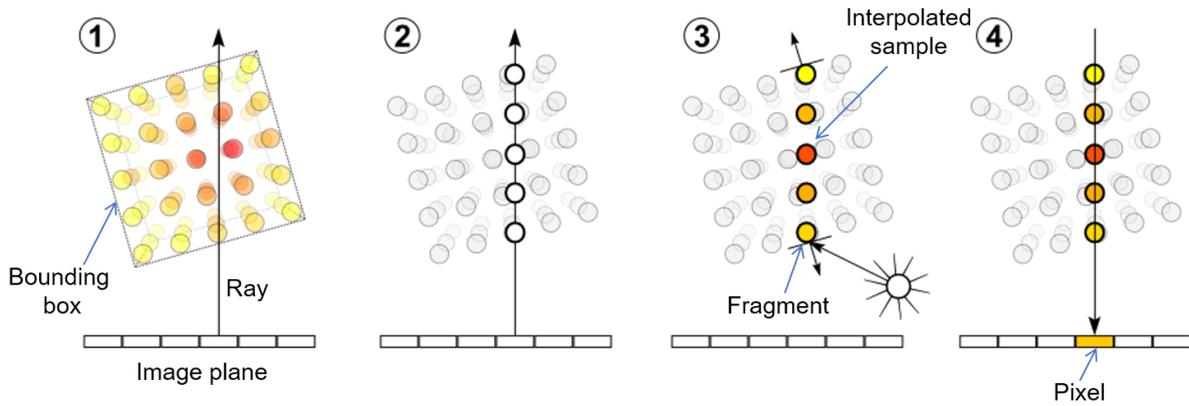


Figure 2.7: In volume ray casting, a single ray is cast into the volume for each pixel that will be used to make up the final image. For demonstration purposes, we show orthographic projection for the rays which make them perpendicular to the image plane. In perspective projection all the rays originate from the location of the viewpoint. The final colour of this fragment on the bounding box is determined by sampling the volume at regular discrete positions along the ray. These samples are then used to evaluate the volume rendering integral [20]. 1. A ray is cast into the volume for each pixel in the image plane. 2. The volume is resampled at equidistant points. 3. The colour at these sample points is determined using some interpolation scheme or a transfer function. 4. The final colour of the pixel in the image plane is determined by compositing the samples along the ray which assigns the final colour to the fragments on the visible faces of the bounding box. Adapted from Thetawave [28].

using the relationship $\kappa(t) := \kappa(s(\vec{x}(t)))$ (Equation 2.3). This leads to the Riemann sum approximation of

$$\tau(0, t) \approx \tilde{\tau}(0, t) = \sum_{i=0}^{\lceil T/\Delta t \rceil} \kappa(i \cdot \Delta t) \Delta t \quad . \quad (2.8)$$

where Δt is defined as the distance between successive resampling locations, and T is the total length of the ray in the volume. Since the optical depth is in the exponent (see Equation 2.5) the summation can immediately be expressed as a multiplication of exponential terms

$$e^{-\tilde{\tau}(0, t)} = \prod_{i=0}^{\lceil T/\Delta t \rceil} e^{-\kappa(i \cdot \Delta t) \Delta t} \quad . \quad (2.9)$$

Now we introduce the *opacity*, usually represented by the Greek capital letter alpha A in alpha blending, by defining

$$A_i = 1 - e^{-\kappa(i \cdot \Delta t) \Delta t} \quad . \quad (2.10)$$

Equation 2.9 then becomes

$$e^{-\tilde{\tau}(0, t)} = \prod_{i=0}^{\lceil T/\Delta t \rceil} (1 - A_i) \quad . \quad (2.11)$$

This allows the opacity A_i to be used as an approximation of the absorption along the i^{th} ray segment, instead of just considering the opacity at a single point along the ray. Similarly, the emission coefficient (Equation 2.3) can be approximated along the i^{th} segment

of the ray as follows:

$$C_i = c(i \cdot \Delta t)\Delta t \quad (2.12)$$

Now that both the emissions and absorptions have been approximated along the ray, with the number of samples along the ray being $n = \lceil T/\Delta t \rceil$, the discretised approximation of the volume rendering integral becomes:

$$\tilde{C} = \sum_{i=0}^n C_i \prod_{j=0}^{i-1} (1 - A_j) \quad . \quad (2.13)$$

Equation 2.13 can be evaluated iteratively while performing *alpha blending* in either a front-to-back or a back-to-front order [29].

Alpha blending

In the context of ray casting, alpha blending refers to the process of compositing the samples along a single ray by blending the colour values based on the opacity (alpha) value.

Equation 2.13 can be written in a recursive form for a front-to-back ordered calculation [29]:

$$C'_i = C'_{i-1} + (1 - A'_{i-1})C_i \quad (2.14)$$

$$A'_i = A'_{i-1} + (1 - A'_{i-1})A_i \quad (2.15)$$

with $i = 1, 2, \dots, n$ and the starting conditions of $C'_0 = 0$ and $A'_0 = 0$.

In a similar way, Equation 2.13 can also be written in a recursive form for a back-to-front ordered calculation:

$$C'_i = C_i + (1 - A_i)C'_{i+1} \quad (2.16)$$

with $i = n - 1, n - 2, \dots, 1$ and starting condition $C'_n = 0$.

Note that for the front-to-back order we need to keep track of the current alpha-value, which is not required for the back-to-front order. However, the major advantage of the front-to-back order is the ability to implement early ray termination, where the iteration terminates once the cumulative alpha value is sufficiently close to 1.0 [29].

Furthermore, note that in the blending equations (Equations 2.14 and 2.16) we use *opacity-weighted colours*, also referred to as *associated colours* [30], where the colours refer to the RGB intensity value of each voxel. An opacity-weighted colour is a colour that has been pre-multiplied with its associated opacity value. This is a convenient notation and is especially useful for interpolation purposes on the graphics hardware. It can be shown that interpolating colour and opacity separately leads to artefacts, whereas interpolating opacity-weighted colours achieves correct results [30]. These artefacts are known as *colour bleeding*, an effect which is illustrated in Figure 2.8.

Implementation

The implementation details of volume ray casting have been relegated to Appendix B. Another variant of volume ray casting known as *pseudo-isosurface* ray casting was also evaluated. With this method, only the surface of similar intensity values, called an isosurface, is visualised. This is similar to what is achieved by popular indirect volume rendering methods such as the marching cubes and marching tetrahedra algorithms [31, 32]. However, unlike indirect volume rendering methods, the ray casting variant does not extract the isosurface as a 3D mesh of polygons. Instead, only the projection of the isosurface onto the bounding box is rendered.

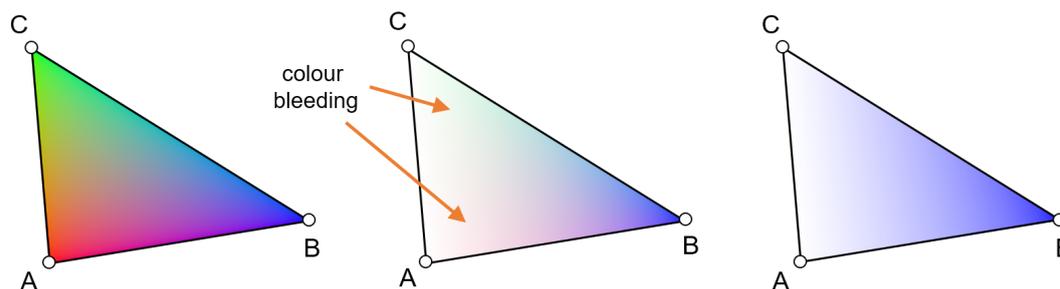


Figure 2.8: An illustration of colour bleeding during interpolation. A triangle is drawn with a different colour specified for each vertex. When vertices A and C are specified to be completely transparent, the colour values in the interior of the triangle are interpolated. Even though the vertices are transparent, their colour is still “bleeding” into the interior of the triangle. In the image on the right, colour bleeding was suppressed by pre-multiplying the vertex colours with their opacity before interpolation. Adapted from Engel *et al.* [20].

2.3.2.3 Texture-based volume rendering

Texture-based volume rendering is an *object-order* approach, which means that the rendering algorithm iterates over arbitrarily generated object slices, known as a *proxy geometry*, and in turn over the fragments on each slice that are visible in the image plane, to produce the visualisation. These slices are then individually projected onto the image plane and composited in the fragment shader using alpha blending. The slices can either be rendered in a front-to-back or a back-to-front order. However, optimisation strategies such as the early ray termination used in ray casting are usually not possible [20].

Since rendering the proxy geometry involves performing repetitive resampling tasks, it can be implemented efficiently using texture mapping on graphics hardware (see Section 2.2.3). For this, the volume data can either be stored in an axis-aligned array of 2D texture images or as a single 3D texture, where each texel (texture element) corresponds to a single voxel. The latter is usually the preferred choice since it allows trilinear interpolation to be used, whereas only bilinear interpolation is supported by the array of texture images. This choice may be limited by the graphics hardware, which may not support 3D textures [20].

Proxy geometry

Proxy geometry refers to the newly generated slices that will be used to visualise the volume instead of the original images in the z-stack constituting the volume data. The volume data is then resampled over the surface of the proxy geometry. This is done in such a way as to improve the resolution of the data visualisation. Specifically, this allows the volume data to be viewed from any direction without looking in-between image slices. This is especially important for fluorescence microscopy images where the number of slice images in the z-stack could be less than 10 and therefore relying on interpolation to create the impression of a 3D structure.

There are several different methods to determine the *proxy geometry*. However, the two methods that are most frequently found in the literature are *object-aligned slices* and *viewer-aligned slices*.

- With *object-aligned slices*, also referred to as *axis-aligned slices* [20], the proxy geometry consists of a stack of planar slices which must be aligned with one of the

major axes of the volume data (x , y or z axis). For this method, the 3D data set is usually stored as an array of 2D textures which is then resampled by the GPU over the surface of the proxy geometry slices [33]. The reason for aligning the slices with a major axis is that two of the 3D texture coordinates are used as actual texture coordinates on the 2D texture while the third coordinate is considered the index of the 2D texture in the array. This third coordinate is therefore not used as a normal texture coordinate. These proxy geometry slices are rendered in back-to-front order and then blended [20].

To allow an interactive rotation of the volume, while avoiding a situation in which the camera looks in-between slices, three arrays of 2D textures (one for each major axis) are required. The texture array on the major axis which minimises the angle between the slice normal and the viewing direction is selected, as is illustrated in Figures 2.9 and 2.10 [20]. Note that, when switching between two different 2D texture arrays, flickering artefacts can occur which diminish the perceived quality of the visualisation. Another drawback of this method is that three arrays of 2D textures need to be stored in memory, which results in a storage requirement three times greater than that required by a single 3D texture. A further disadvantage is that the sampling rate is not consistent between the slices when viewing the volume from different angles as is illustrated in Figure 2.10B. This leads to intensity variations as the camera moves around the volume, which could be interpreted as variations in density, and consequently to inaccuracies in the perceived opacity. This is especially problematic with fluorescence microscopy z-stacks that contain very few images in the z-stack. Finally, this method allows only the application of bilinear interpolation, which can lead to reduced rendering quality compared to trilinear interpolation.

A chief advantage of using object-aligned slices is that it is less computationally intensive to calculate the geometry of the slices, an operation that is usually performed on the CPU, than the viewer-aligned slices approach. It is, therefore, useful when performance should be optimised, which made it the historically preferred method for real-time volume rendering. Furthermore, since this approach uses an array of 2D textures, it can be implemented on older GPUs which do not support 3D textures.

- The *viewer-aligned slices* approach addresses several of the problems associated with the use of 2D texture arrays by using 3D textures. While bilinear interpolation must be applied to arrays of 2D textures modern graphics hardware is capable of performing trilinear interpolation over the interior of the proxy geometry polygons when 3D textures are used. In this way, the proxy slices no longer need to be bound to major axes of the volume data. The inconsistent sampling rate for different axes that is associated with object-aligned slices is also sidestepped by using arbitrarily positioned slice polygons in the 3D texture domain, since we can now ensure that the same number of slices are always generated and that they are parallel to the image plane instead of to one of the object axes. By using proxy slices that face the view direction, this approach closely mimics the sampling used by the ray-casting algorithm [20]. Notice, however, that in the orthogonal (parallel) projection case there is a consistent sampling rate for all viewing rays (Figure 2.11a), while for the perspective projection the sampling rate is still not consistent (Figure 2.11b). However, according to Engel *et al.* [20] such effects are only noticeable when the

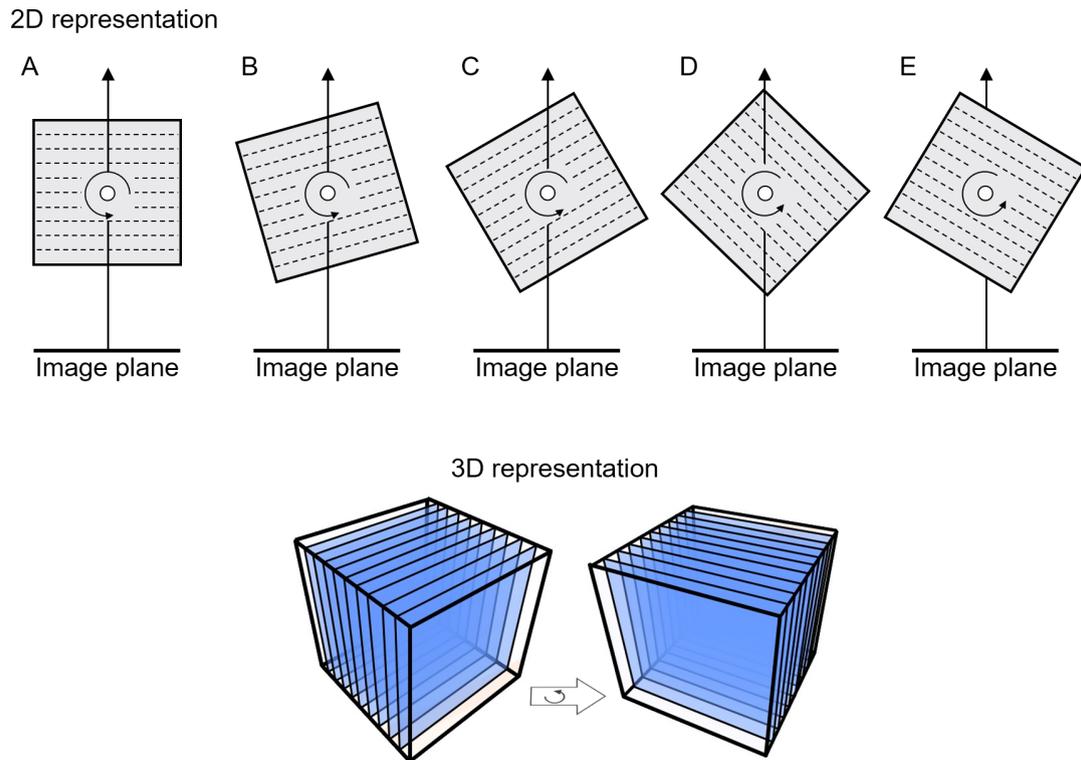


Figure 2.9: Switching between two texture arrays in the case of object-aligned slices in order to minimise the angles between the viewing direction and slice normals. Between image (C) and (D) the 2D texture arrays were switched. Adapted from Engel *et al.* [20].

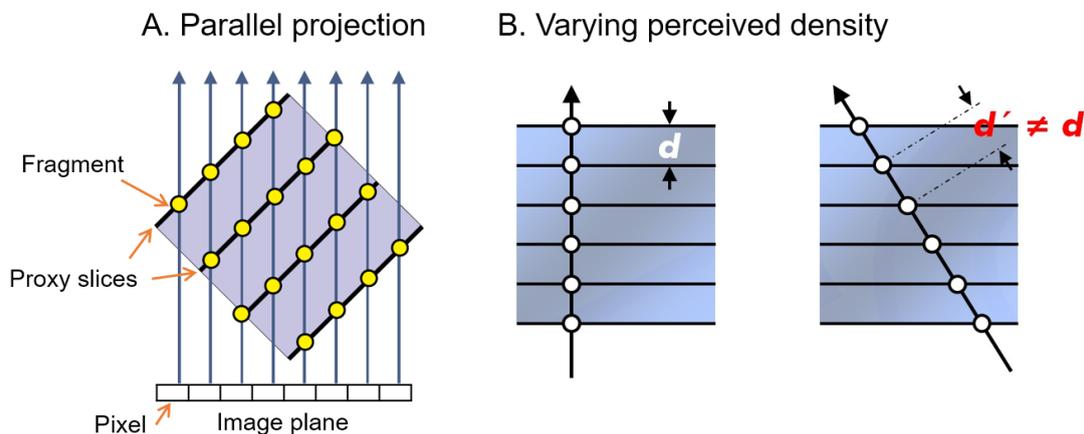


Figure 2.10: Object-aligned texture-based volume rendering. A. A parallel projection of samples (yellow) on the object-aligned slices. B. Demonstration of how the distance between adjacent slices depend on the viewing direction which affects the perceived volume density. Adapted from Engel *et al.* [26].

field-of-view is extremely large.

The chief disadvantage of viewer-aligned slices is that trilinear interpolation is significantly more computationally expensive and hence more time consuming than bilinear interpolation. Furthermore, the proxy geometry must be recalculated, either on the CPU or in a vertex shader, every time the viewing direction changes. For these reasons, more powerful hardware is required to achieve acceptable frame

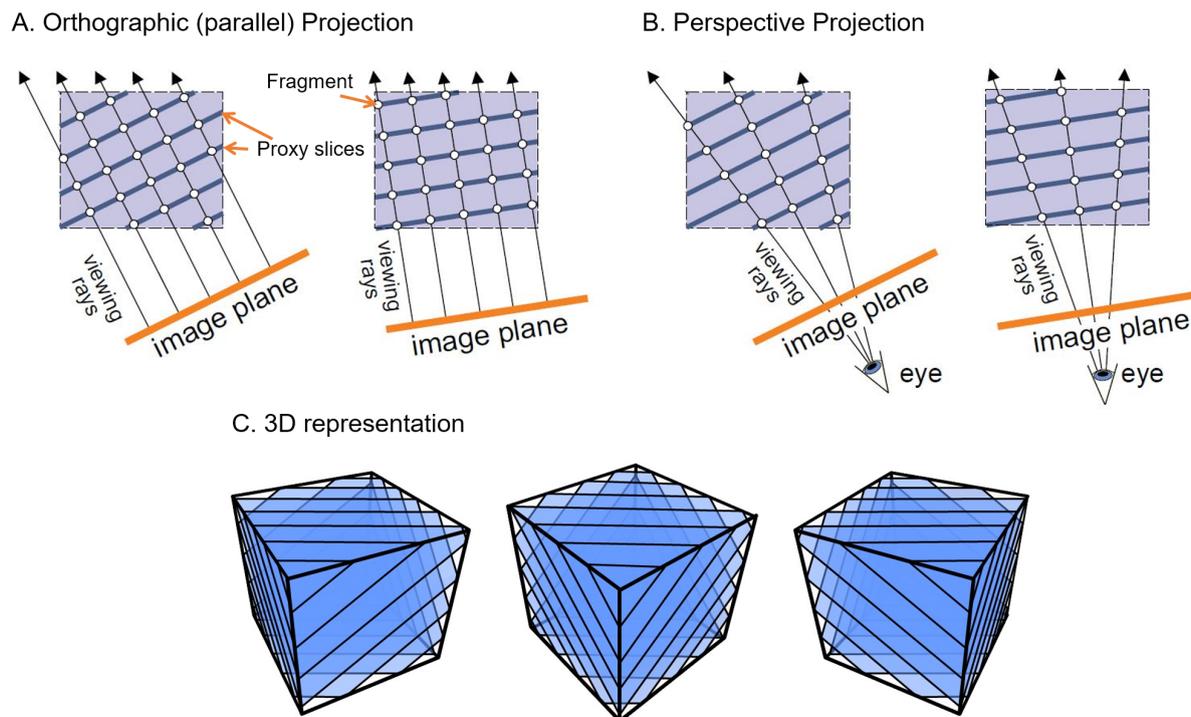


Figure 2.11: Sampling points illustrated for viewer-aligned slicing for orthographic (or parallel) projection (A) and perspective projection (B) respectively. These sampling points are the fragments on the proxy slices which are blended together to create the final image on the image plane. C. Demonstration of how the proxy geometry slices changes as the viewing angle change. Adapted from Engel *et al.* [26].

rates. Furthermore, a large amount of texture memory must be available to store the 3D texture in the case of large or high-resolution volume data [20]. The total required memory is, however, less than what is necessary for object-aligned slices approach since only one copy of the texture data needs to be stored.

The implementation details for viewer-aligned texturing slicing have been relegated to Appendix A.

2.3.2.4 Other direct volume rendering methods

There are several other direct volume rendering methods which we also considered during the development of our visualisation system. Two notable approaches are the *shear-warp volume rendering*, which is similar to the object-aligned textured slicing approach that we discussed in Section 2.3.2.3, and *splatting* [20].

These approaches, however, prioritise performance over quality. Since our objective is the accurate visualisation of biological data to support research and scientific discovery, quality is the priority. Furthermore, due to recent advances in GPU performance, we were able to produce real-time volume rendering even for the most computationally intensive methods, specifically viewer-aligned texture slicing.

2.3.2.5 Comparison of volume rendering techniques

Even though both texture-based volume rendering and volume ray casting can produce similar results, there are differences which justify two implementations.

Texture-based volume rendering allows the user to move into the volume and still see the visualisation since the proxy geometry is rendered as slices throughout the 3D volume. This is not possible with volume ray casting since the rendering is performed on the faces of the bounding box. Hence the visualisation disappears as soon as the user moves inside the bounding box. By allowing visualisation from within the rendered volume, the texture-based approach allows a more immersive sample investigation, especially when visualised in virtual reality.

Volume ray casting, on the other hand, generally produces a sharper rendered image, since texture-slicing can lead to inconsistent sampling rates as we illustrated in Figure 2.10B. Furthermore, it allows pseudo-isosurface rendering, which makes a more detailed investigation of sample regions with similar intensity values possible. This is, for example, common at cell boundaries.

We used implementations of both volume ray casting as well as the viewer-aligned texture slicing for our visualisation system.

2.4 Virtual reality

Due to their complex nature, a good understanding of fluorescence microscopy samples cannot be achieved by visualising the 3D reconstructed sample on a 2D display, such as a computer screen. This is due to the loss of depth information in the projection which can sometimes lead to ambiguity in the visualisation. Recent developments in virtual reality, make it possible to produce an unambiguous visualisation of a sample on a 3D display using readily available consumer hardware. With this hardware, a different image is presented to each eye, providing a truly 3D visualisation. New technology also provides more intuitive ways for the researcher to interact with the sample, an aspect that will be further investigated in Chapter 3.

Virtual reality refers to the simulated presence of a user in a computer-generated 3D environment. The user can explore and interact with virtual objects while immersed in this environment. Most current implementations of virtual reality (VR) make use of a stereoscopic *head-mounted display* (HMD)². There is usually also some form of input device which allows the user to interact with the environment in a natural way. This can range from a standard keyboard and mouse to devices that were specifically designed for virtual reality, such as wired gloves or omnidirectional treadmills. Input devices may also provide haptic feedback, which is tactile information about the current interaction in the VR scene.

2.4.0.1 Stereo rendering concepts

In the case of VR, two virtual cameras are set up, one for each eye. The scene must therefore be rendered twice, once for each camera, and each rendered image displayed only to the corresponding eye. This is known as *stereo rendering*. The two virtual cameras should be separated by the same distance in virtual space as the user's pupils are separated. This separation is known as the *interpupillary distance* (IPD) and is on average 61.1 ± 3.5 mm for women and 63.6 ± 3.9 mm for men [34]. Setting up these two virtual cameras corresponds to a translation, but not a rotation. It is the parallax effect that accompanies the translation that leads to the stereoscopic effect (Figure 2.12).

²We used the Oculus Rift for this project, however, our visualisation system is independent of the choice of HMD.

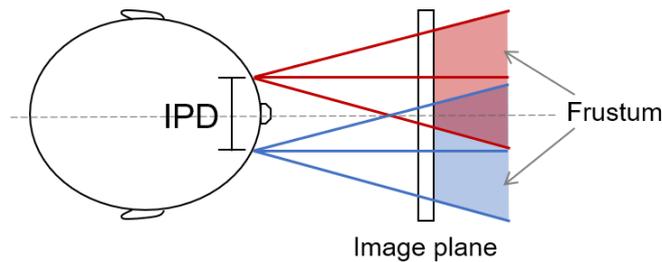


Figure 2.12: Example of how the IPD relates to the rendering image plane.

2.4.0.2 Other VR related factors

When rendering a virtual environment in VR, three additional factors must be taken into consideration, namely the field-of-view (FOV), distortion correction and chromatic aberration correction. If these are neglected, the VR environment may appear warped, which can induce simulator sickness. Simulator sickness refers to symptoms of discomfort, such as nausea, that arise from simulated environments due to a conflict between visual, vestibular (balance), and proprioceptive (bodily position) senses [35].

The FOV is the extent of the observable world that is seen at any given moment. It is important to disambiguate two different usages of FOV. The one is *display field-of-view* (dFOV), which refers to the part of the user's physical visual field which is occupied with VR content. The other is *camera field-of-view* (cFOV), which refers to the range of the virtual world can be seen through the rendering cameras. To prevent simulator sickness, it is important that cFOV and dFOV match exactly when using VR.

The lenses that are used in VR HMDs magnify the image on the screen in such a way as to provide a wide FOV to the user, in order to enhance the sense of immersion. This, however, can cause the image to be significantly distorted. If this distortion is not compensated for, the user would see the images with *pincushion distortion* (Figure 2.13). To counteract this distortion, the software must perform post-processing on the rendered images with an equal and opposite *barrel distortion* (Figure 2.13) in order to cancel out the pincushion distortion.

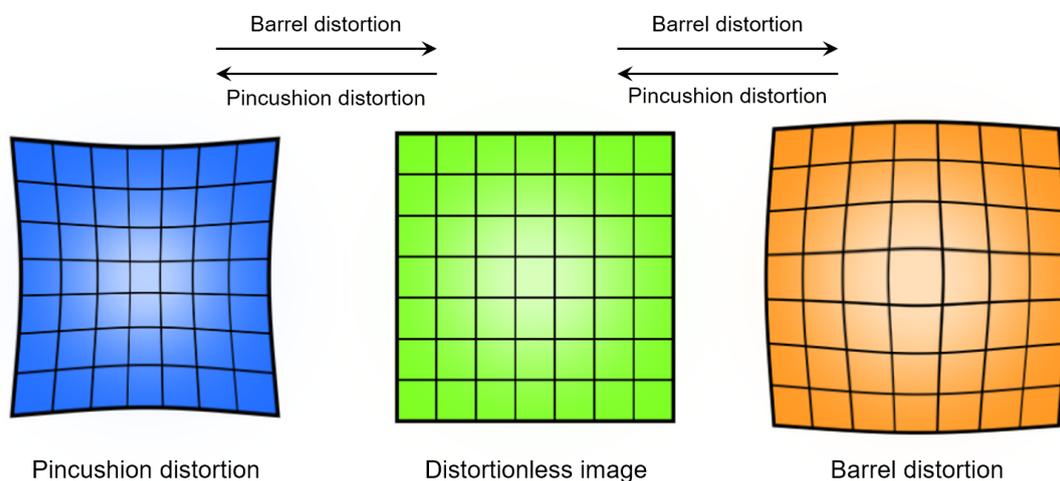


Figure 2.13: Illustration of how distortion affects straight lines. Opposite distortions can be used to rectify a distorted image. Adapted from Kleciak [36].

Finally, the lenses in an HMD can introduce *chromatic aberration*. This distortion causes a colour separation effect at the edges of the lens. The effect is due to the refractive index of the lens that varies for different wavelengths of light. Chromatic aberration can be corrected by pre-transforming the image so that, after it has been refracted by the lens, it looks normal.

2.5 Conclusion

In this chapter, we have presented an overview of the main concepts that we employed to develop the visualisation system on which the following chapters will be based. We discussed how fluorescence microscopy z-stacks can be three-dimensionally reconstructed using direct volume rendering. The two main approaches that we considered were volume ray casting and texture-based volume rendering. To achieve real-time interactive volume rendering, these algorithms leverage their inherently parallel nature and the processing power of current graphics processing units (GPUs). To this end, a brief overview of graphics programming, as well as some associated concepts, was provided. Finally, we discussed the benefit that virtual reality (VR) provides for the visualisation and analysing fluorescence microscopy samples. This will be further investigated in the following chapters.

In the next chapter, we will expand our visualisation system and specifically apply it to the analysis of colocalisation in fluorescence microscopy. Since modern VR has not previously been applied to the visualisation of these samples, we will also investigate different user interface methods to determine which provide the best productivity and user experience.

Chapter 3

Virtual reality assisted microscopy data visualisation and colocalisation analysis

The content of this chapter was published as the following research article:

Theart, R.P., Loos, B. and Niesler, T.R., 2017. Virtual reality assisted microscopy data visualization and colocalization analysis. *BMC Bioinformatics*, 18(2), p.64.

In the previous chapter we developed a virtual reality based system that uses direct volume rendering (DVR) to three-dimensionally reconstruct fluorescence microscopy image stacks in real-time and visualise the result using VR. In this chapter we expand this system and apply it to the task of colocalisation analysis, which is an important technique used in biological and physiological studies to investigate fluorescence microscopy images.

3.1 Introduction

In the past, biological visualisation has been limited to rendering on two-dimensional displays. The use of three-dimensional displays for quantitative signal assessment, including the precise signal selection, the processing of multiple signals and the determination of their spatial relationship to one another, has received very little attention.

Recent developments in virtual reality (VR) have led to lightweight headsets that offer high resolution, low latency head tracking, and a large field-of-view. These advances make VR an attractive technology to use in biological visualisation. It allows immersive three-dimensional visualisation, as opposed to a three-dimensional rendering on a two-dimensional display [2]. Because the visualisations are based on a true three-dimensional awareness of the sample, they offer an unambiguous representation as well as a more intuitive process of interaction. This can aid the process of scientific investigation and discovery.

One important factor that has limited the usefulness of microscopy volume visualisation in the past was the low rendering speed. Just seven years ago it was challenging to exceed 20 frames per second (10 frames per second per eye) on consumer computing equipment [2, 3]. This is far from adequate for VR, which requires a frame rate of at least 60 frames per second per eye in order to achieve an immersive experience [4]. If this frame rate is not maintained, simulation sickness may result [35]. With recent improvements in consumer graphics performance, the frame rates necessary for VR can be achieved with reasonable rendering quality.

When using VR for visualisation, traditional interaction methods such as a keyboard and mouse may no longer be appropriate. We, therefore, consider two alternatives¹: the use of a gamepad, in combination with the VR headset's head tracking, and the use of hand tracking for gesture control.

To assess the quality, suitability and characteristics of our VR system, we applied to the analysis of colocalisation in a fluorescence-based three-dimensional biological sample acquired from a confocal microscope. In fluorescence microscopy, colocalisation refers to the codistribution of two fluorescent probes. This can be used to determine whether two molecular structures associate one another. Colocalisation can occur in two ways: co-occurrence and correlation. *Co-occurrence* is the simple spatial overlap of two probes, which are not necessarily proportional. *Correlation* refers to two probes that not only overlap but also codistribute in proportion to one another within and between structures [37]. There are several metrics that can determine whether two probes, or fluorescence channels, merely co-occur or whether they are also correlated. We implemented a selection of such metrics that are commonly used in colocalisation analysis.

3.1.1 Colocalisation metrics

In order to determine co-occurrence, both *Manders' colocalisation coefficient* (MCC) and the *percentage colocalisation* were calculated. For both co-occurrence and correlation, *Pearson's correlation coefficient* (PCC) and *Manders' overlap coefficient* (MOC) were calculated [37]. All of these metrics are calculated above an intensity threshold that is manually determined during the acquisition of the sample images.

3.1.1.1 Percentage colocalisation

Percentage colocalisation, or overlap, is a metric that reflects only the co-occurrence of the two fluorescence probes. It is useful when trying to quantify the extent to which one channel colocalises with another channel. The percentage of colocalised voxels are calculated as follows:

$$x\% = \frac{c}{n} \quad (3.1)$$

In Equation 3.1, x is the percentage colocalisation, c the total number of overlapping voxels between the two fluorescence channels and n is the combined total number of voxels for both channels. Only the voxels for which the fluorescence intensity is above the corresponding channel intensity threshold are counted for these totals. It is often useful to compare this metric with the Manders' colocalisation coefficients.

3.1.1.2 Pearson's correlation coefficient

Pearson's correlation coefficient (PCC) is a measure of the strength of linear association between two variables [38], and can be calculated using the following equation:

$$\text{PCC} = \frac{\sum_i (R_i - \bar{R}) \times (G_i - \bar{G})}{\sqrt{\sum_i (R_i - \bar{R})^2 \times \sum_i (G_i - \bar{G})^2}} \quad (3.2)$$

¹At the start of the project (2015), the Oculus Touch motion controller system was not yet released, and for that reason, we did not consider it as a possible input interface when developing our visualisation system.

In Equation 3.2, R_i and G_i refer to the respective intensity values of the first and second colour channels for voxel i , respectively. \bar{R} and \bar{G} refer to the corresponding mean intensities calculated across the entire sample volume. All these values are calculated for the voxels that are above the corresponding channel intensity threshold [37]

PCC values lie in the range -1 to +1 for perfectly inverse correlated and perfectly correlated fluorescence channels, respectively. A PCC value near zero reflects two uncorrelated fluorescence channels.

3.1.1.3 Manders' overlap coefficient

Manders' overlap coefficient (MOC) is a similar but alternative metric to the PCC and can be described by the following equation:

$$\text{MOC} = \frac{\sum_i (R_i \times G_i)}{\sqrt{\sum_i R_i^2 \times G_i^2}} \quad (3.3)$$

Unlike PCC, R_i and G_i are not subject to mean normalisation in the calculation of the MOC and therefore this quantity cannot become negative. In some cases, this aids in the interpretation of these values. An unfortunate consequence of this, however, is that that even uncorrelated channels could produce a high correlation value. This reflects the fact that MOC is more independent of signal proportionality, and is mainly sensitive to co-occurrence [37].

3.1.1.4 Manders' colocalisation coefficient

When it is important to determine what fraction of one channel colocalises with another, the Manders' colocalisation coefficient (MCC) is a metric that is widely used in biological microscopy. Since it measures the fractional colocalisation it provides a more intuitive and direct metric of examining colocalisation [37]. It can be calculated as follows:

$$\text{M1} = \frac{\sum_i R_{i,colocal}}{\sum_i R_i} \quad (3.4)$$

where $R_{i,colocal} = \begin{cases} R_i, & \text{if } G_i > 0 \\ 0, & \text{if } G_i = 0 \end{cases}$

$$\text{M2} = \frac{\sum_i G_{i,colocal}}{\sum_i G_i} \quad (3.5)$$

where $G_{i,colocal} = \begin{cases} G_i, & \text{if } R_i > 0 \\ 0, & \text{if } R_i = 0 \end{cases}$

Measuring MCC is more complicated, because the input values can rarely be taken directly from the original image. That is because light leakage is often present in the system, and these "background" pixels must first be removed. These "background" pixels can be removed by a threshold value, but unfortunately MCC is very sensitive to the value that is chosen for the threshold [37]. In our case a low threshold, as mentioned in the percentage colocalisation section, was used to remove the background.

3.1.2 Region of interest

The region of interest (ROI) is a selected sub-volume of the complete microscopy volume sample. It is a specific part of the sample that the biologist would like to investigate while ignoring the rest. For example, the amount of colocalisation in the nucleus of a cell may be of particular interest. For colocalisation metrics to yield meaningful information, it is generally very important to carefully outline the region of interest.

3.1.3 Methods to visualise colocalisation

Instead of simply calculating colocalisation, it is often more intuitively useful to visualise it. Two techniques are commonly used to achieve this. Firstly, a spatial sense of the colocalisation can be obtained by replacing, or *superimposing*, the colocalised voxels with white voxels. This method is, for example, useful for identifying regions of a cell or compartments where certain molecules colocalise [37]. Secondly, the outputs of the two colour channels can be represented as a *scatter plot*, where the intensity of one channel is plotted against the intensity of the other, either using the original display colours or by representing the frequency of overlap between the colour intensities by pseudo-colours in the scatter plot, which aids interpretation [39]. Scatter plots allow one to determine visually how well the two colour channels are correlated [37].

In Chapter 5 we develop a new method to visualise the correlation distribution of the colocalisation between the fluorescence channels throughout the sample by using a colour map.

3.1.4 Virtual reality

Virtual reality (VR) refers to a 3D computer generated environment into which the user is immersed. The user can explore and interact with the virtual objects in this environment. Modern VR is mostly implemented using a stereoscopic head-mounted display (HMD), such as the Oculus Rift [40]. The HMD presents a separate image to each eye in order to achieve 3D perception of the scene. In the case of the Oculus Rift, head position tracking is accomplished via an infrared cameras, and head orientation tracking is achieved using a 3-axis gyroscope, an accelerometer and a magnetometer. This allows the environment to be updated in response to head movements and allows for an immersive experience. This immersion is further enhanced when the HMD is used in conjunction with a hand tracking device such as the Leap Motion [41], or a physical haptic input device, such as a gamepad.

3.1.5 Previous work

Very little work has been published in the field of immersive microscopy visualisation. In the following we briefly describe four relevant systems. As far as we are aware, no system has been used to perform colocalisation analysis in a VR environment.

CAVE2 Hybrid Reality Environment

CAVE was first developed in 1992 and upgraded to CAVE2 in 2012 at the Electronic visualisation Laboratory, University of Illinois at Chicago. It is capable of 3D visualisation in general, but has been extensively applied to scientific visualisation.

CAVE2 is a cylindrical system, 7.3 meters in diameter and 2.4 meters high. It consists of 72 near-seamless, off-axis-optimised passive stereo LCD panels, creating an approximately 320 degree panoramic environment for displaying information at 37 megapixels (in stereoscopic 3D) or 74 megapixels in 2D and at a horizontal visual acuity of 20/20. Users can interact with visualised objects using a tracked wand, called a navigational controller. For 3D visualisations the user wears stereoscopic glasses. The position and orientation of the user's head and the navigational controllers are tracked with an array of 10 infrared cameras [42].

SkinExplorer

The SkinExplorer is a VR platform for visualising 3D reconstructed confocal microscopy images, specifically for the investigation of skin structures [5]. The system was designed to use either a large-screen 3D display, where interactions are performed using an ART Flystick3 in combination with an ART SMARTTRACK tracking system, or a desktop 3D display, where interactions is achieved using a Microsoft Kinect in combination with a joystick or a mouse for navigation. In both cases the user wears stereoscopic glasses. Head movement is tracked using sensors attached to the glasses, and in the desktop approach the Kinect automatically tracks the face to determine the position and orientation of the eyes. A GUI displays 3D widgets alongside the visualisation as well as running on a remote touch interactive device [5].

BRAINtrinsic

BRAINtrinsic is a web-based 3D visual analytics tool that allows users to intuitively and iteratively interact with connectome data of the human brain. The system implements a visualisation platform that reconstructs connectome's intrinsic geometry, which is the topological space as informed by brain connectivity. BRAINtrinsic was developed with VR in mind and is fully compatible with the Oculus Rift [6].

3D+Time Brain View

The 3D+Time Brain View system visualises functional Magnetic Resonance Imaging (fMRI) data gathered from users exposed to unfamiliar spoken languages. The system illustrates the temporal evolution of participants' brain activity as they are introduced to a foreign language by displaying these clusters as they change over time. This is achieved by reconstructing the fMRI data using volume ray casting and displaying it on a projected stereoscopic display, with the user wearing polarised glasses. According to the developers of this system, they have been experimenting with off-the-shelf interaction technologies such as the Microsoft Kinect and Leap Motion, although no details about their implementation could be found [43].

Relation to proposed system

In contrast to the CAVE2 system, our algorithms are implemented on hardware that is fairly easy to obtain and much less costly. It is therefore accessible to researchers. It is also much more portable and does not require the constriction of specialised hardware. The SkinExplorer is interesting because it has also been applied to confocal microscopy. Like the CAVE systems, it is based around stationary 3D displays. The VR headset

approach we propose offers better 3D immersion. Furthermore, we allow the user's hands to be used to manipulate the sample without the need for hand held devices.

Like the system we propose, the BRAINtrinsic system makes use of modern VR to aid analysis. However, instead of volumetric reconstruction of fMRI images, the system reconstructs connectome data. In contrast, the 3D+Time Brain View system uses volume rendering to reconstruct the fMRI images of the brain. Also in common with our system, it incorporates hand tracking. However, no information on the implementation or performance thereof could be found.

3.2 Methods

We set out to create a system that would allow improved confocal microscopy data visualisation and colocalisation analysis. To achieve this goal we worked closely with domain experts to identify limitations in their current analysis capabilities as well as to verify that our prototypes were indeed functional. The main limitation identified in current systems is the lack of flexibility in 3D sample navigation and interrogation, and to dissociate a region of interest from the main sample which is required to perform selective analysis. A second aspect is inability to threshold the fluorescence channel intensities, and to provide an indication of how they affect colocalisation in 3D by displaying it that accordingly. These aspects will be addressed by the adapting the system developed in Chapter 1.

A main objective was to make the colocalisation analysis process more intuitive. In this respect, virtual reality was a very attractive option, since it offers immersion into a three-dimensional environment that is already familiar to any user. The fact that the samples could be visualised in 3D is an important advantage over currently prevalent methods, where the sample visualisation is displayed on a two-dimensional screen, and where relative spacial positions are difficult to judge due to the lack of depth.

The user can choose to visualise the microscopy sample using either the texture-based or the volume ray casting rendering methods (see Section 2.3.2), depending on what they want to accomplish.

In order to allow colocalisation to be calculated in an interactive way, a graphical user interface (GUI) was implemented which allows the user to manipulate the rendering parameters of the sample. These parameters include the noise filtering threshold, the global opacity and the opacity of the different colour channels. Further GUI panels allow more detailed colocalisation analysis as well as region of interest selections. Examples of these GUI panels are shown in Figures 3.1 and 3.2.

Our system was implemented using the Unity engine [14]. Unity uses a left-hand coordinate system, with the x-axis pointing towards the right of the view, the y-axis pointing up and the z-axis pointing away from the user. This is the coordinate system that is used for scaling the ROI. Unity natively supports the use of VR headsets. It automatically creates two rendering cameras, one for each eye, and updates the position and orientation of these cameras in the 3D scene based on the position and orientation tracking sensor data.

3.2.1 Sample preparation

In order to test our visualisation system, mouse embryonic fibroblast (MEF) cells were stained with anti-DNA, anti tubulin and anti-actin-Alexa-633 antibody, followed by incu-

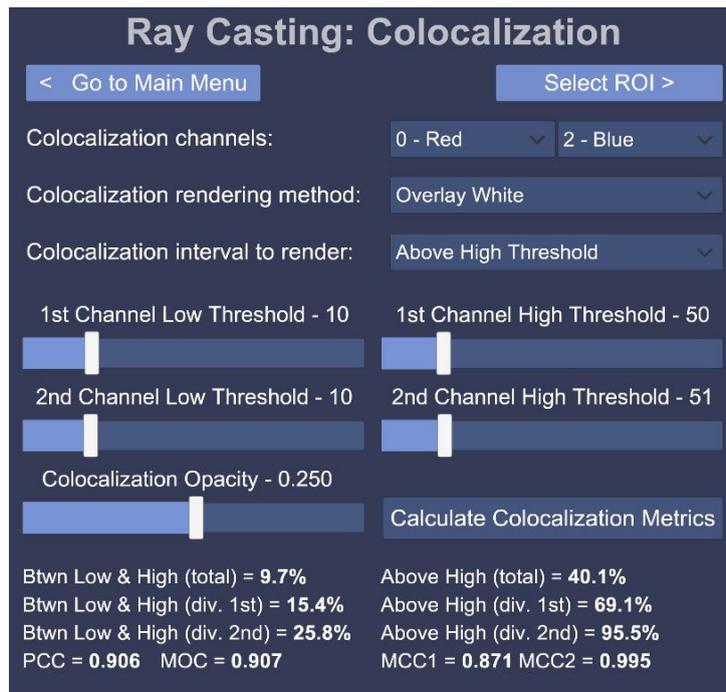


Figure 3.1: An example graphical user interface (GUI) presented by our system. Here the colocalisation settings GUI is shown. In this case the colocalisation parameters can be adjusted, such as the channels that should be checked for overlap, how the rendering should be performed, what the high and low threshold should be as well as the opacity of the colocalised voxels.

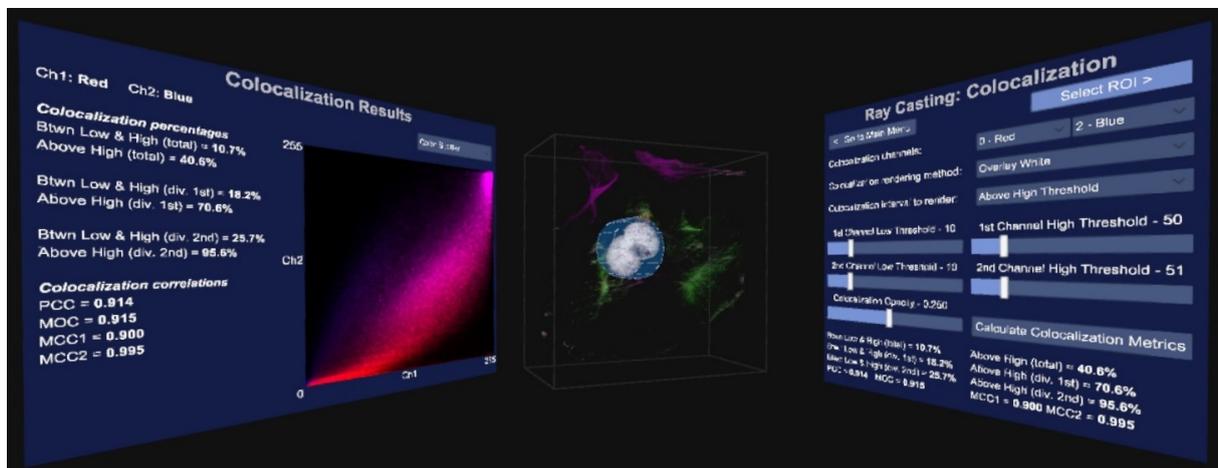


Figure 3.2: The broader GUI view. The user has GUI panels on both sides that directly face the user when looking in that direction. The colocalisation GUI is to the right, the volume sample in the centre and the results of the colocalisation analysis to the left. The volume sample and results are updated immediately as the parameters in the GUI are varied.

bation with conjugated Alexa-488 donkey-anti-rabbit and Alexa-568 donkey-anti-mouse secondary antibody (Life Technologies). Nuclei were counterstained using the DNA intercalating fluorochrome Hoechst 33342. Image acquisition was performed using the Zeiss LSM 780 confocal microscope equipped with a GaAsp detector and images were acquired through z-stack acquisition, with an increment of $\pm 0.4 \mu\text{m}$ between image frames. The

AxioCam MRm camera was utilized to capture images.

The sample can be seen in Figure 3.3d which shows a mammalian cell with 2 nuclei (blue) surrounded by small mitochondrial DNA fragments (red) as well as a thin microtubule network (green) and a thick structural actin network (magenta). Note that due to the blue and red channels overlapping, the nucleus also appears magenta. The tubulin and actin network facilitate the protein transport as well as the cell stability, creating a cellular skeleton.

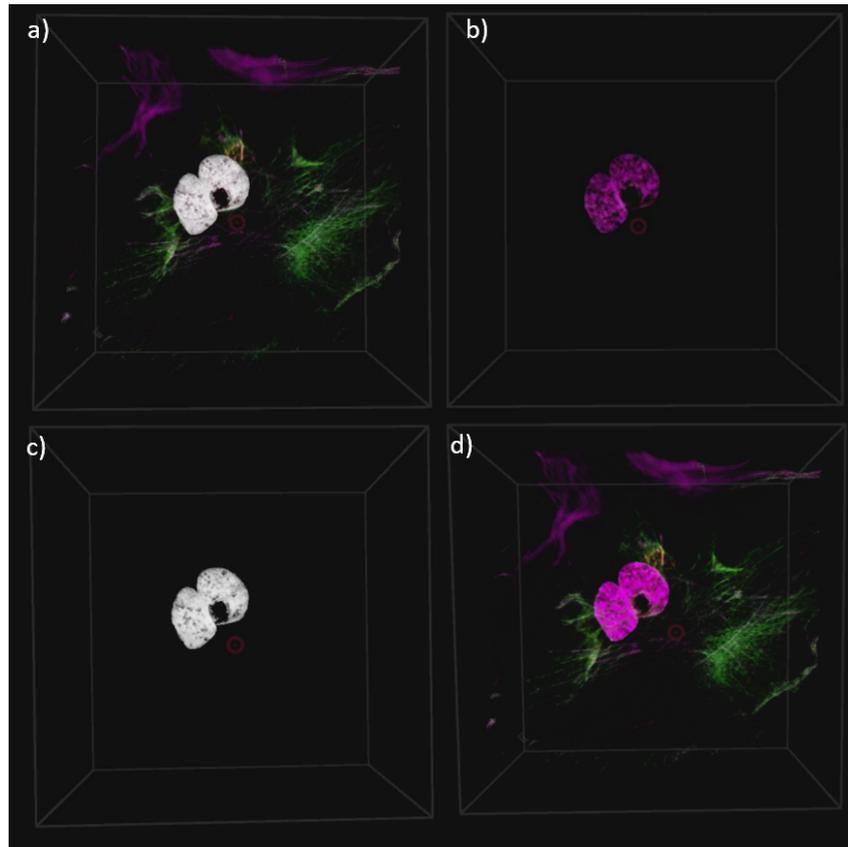


Figure 3.3: The colocalisation can be rendered in four different ways: a) Overlaying the colocalised voxels as white voxels on the original volume sample. b) Only showing the colocalised voxels in their original fluorescence colours. c) Only showing the colocalised voxels as white. d) Not showing any colocalisation. Note that due to the blue (nuclei) and red (mitochondrial DNA) channels overlapping, the nucleus also appears magenta.

3.2.2 Input interfaces

When a user wears a VR headset they are visually cut off from the world. This makes traditional input such as a keyboard, mouse, pen or touch interface impractical. Two alternative methods for interacting with the system were therefore implemented. The first uses the Leap Motion hand tracking system [41], which allows the user's hands to be visible inside the VR environment. The user can therefore use his or her hands to interact with the GUI and the sample itself in an already familiar and intuitive way. While this input method has the important advantage of absolute intuitiveness, due to inaccuracies in the hand tracking process it can however lead to some frustration when

precise interactions are required. This is compounded by the lack of haptic feedback that would be present when touching real objects.

As a second input method we combined the VR headset's built-in head tracking with a traditional gamepad. Since a gamepad is a physical device that the user can hold, it allows for finer control over the visualisation. Button presses are combined with the direction of the user's gaze, as determined from the HMD's head tracking system. Most gamepads also provide haptic feedback, which gives the user a greater sense of physical interaction. However, a gamepad is less intuitive and requires the user to learn which input interactions are mapped to which functions. Also, since it is not part of the rendered scene, it is less immersive.

Hand tracking

Hand tracking was achieved using the Leap Motion [41], which uses a stereoscopic infrared camera, along with a Unity supported software API, to interpret the camera data. When attached to the front of the VR headset, the user's hands can be tracked and rendered inside the virtual environment. Our system allows the user to translate, scale and rotate the sample using intuitive hand gestures, such as pulling and turning.

Since hand tracking is sometimes inaccurate, these gestures are designed to be simple and unambiguous. In designing suitable gestures, we have been guided by other standard VR gesture based applications designed around the Leap Motion as well as Leap Motion's "VR Best Practices Guidelines" [44, 45, 46]. For example, we do not rely on small finger movements.

Since there is no physical feedback from the virtual content, visual feedback is used to inform the user when gestures or interactions with virtual elements are detected. We have accomplished this by rendering a bounding box around the volume sample and changing its colour, based on the currently detected interaction. This was paired with optional text hints that were momentarily displayed over the sample.

In order to translate, scale or rotate the sample, the user must move both hands to be partially or totally inside the bounding box. Both hands must then perform the pinch gesture (index finger touches thumb). When this double pinch is detected, translation can be performed by moving both hands in the same direction simultaneously, while scaling can be performed by pulling the hands apart or pushing them closer together. The sample is rotated when the hands rotate relative to each other. Both rotation and translation can also be performed by pinching with only one hand, and moving the hand around for translation and turning the hand for rotation. We found that this allows for more precise rotations.

The user can interact with the GUI by touching the GUI elements in 3D space, in a similar way that a touch screen would be used.

Inaccuracies of the hand tracking system

The hand tracking system provided by the Leap Motion and the Orion SDK is currently the best supported consumer tracking system that also integrates easily with a virtual reality headset. It also provides the most stable hand and finger tracking that we are aware of. There are, however, several technical limitations to the system that can cause some difficulty for novice users.

Firstly, due to the use of pinch gestures for interacting with the sample, precise finger tracking is necessary. However, when the user's hands move too far away from the tracking

device, finger tracking sometimes becomes unreliable. This also makes interaction with GUI elements using a single finger difficult. These problems can, however, be mitigated by requiring the user to move physically closer to the element in question.

Secondly, because the hand tracking system needs line-of-sight to accurately track each finger it becomes inaccurate when the fingers are obstructed by, for example, the other hand [46].

Unfortunately, neither of these problem scenarios are instinctively clear to novice users. Even when they are cautioned not to perform such problematic hand movements, the immersive nature of virtual reality makes the adherence to this advice difficult. The frequency with which these difficulties were experienced varied greatly from user to user, and it is expected that they would be largely overcome as users become experienced in using the interface.

Gamepad

We made use of a traditional gamepad with two analogue sticks and four front facing buttons, a directional pad and four trigger buttons. Translation, scaling and rotation were performed using the analogue sticks in conjunction with button presses. We combined the gamepad input with the head tracking provided by the VR headset. Using the head tracking, a 3D cursor is rendered in the centre of the display. When the 3D cursor hovers over a GUI element, it is highlighted and raised slightly to indicate to the user which element they will interact with when a button is pressed on the gamepad.

Interaction with the GUI elements is accomplished by moving the 3D cursor over the element in question and pressing a button on the gamepad. When sliders are selected in the GUI their value can be changed using the direction pad. This offers an advantage over the hand tracking, since the user can change the rendering parameters without having to continue looking the GUI.

3.2.3 Colocalisation visualisation and analysis

The GUI layout for colocalisation analysis, which shows several colocalisation metrics as well as scatter plots, can be seen in Figures 3.1 and 3.2. Figure 3.2 shows that the GUI panels are angled to face the user when looking in that direction, with the volume sample in the centre. This setup increases the immersion in the virtual environment. Colocalisation is visualised in real-time as the settings are changed.

Figure 3.1 shows how the GUI allows the user to select the two channels that should be considered for colocalisation analysis. In order to assist with the analysis the user can select to overlay the colocalised voxels on the volume sample as white voxels, or to render only the colocalised voxels either in their original colours or in white. These rendering options are illustrated in Figure 3.3. Furthermore, the thresholds used in the colocalisation metric calculations as well as the rendering opacity of the colocalised voxels can be adjusted. Once these parameters have been optimised interactively, the colocalisation metrics and scatter plots can be calculated. All these are only calculated within a pre-selected *region of interest* (ROI), which is discussed in the next section.

3.2.4 Selection of the region of interest (ROI)

In order to effectively investigate the colocalisation between two colour channels, a good region of interest (ROI) selection tool is required. Three different ROI selecting tools

were implemented, namely the *box*, *cylinder* and *freehand* tools. Example selections with these tools can be seen in Figure 3.4. When the user is in the ROI selection mode, the same interactions used to manipulate the sample are used to manipulate the ROI. The user additionally has the ability to scale the selection along a specific axis.

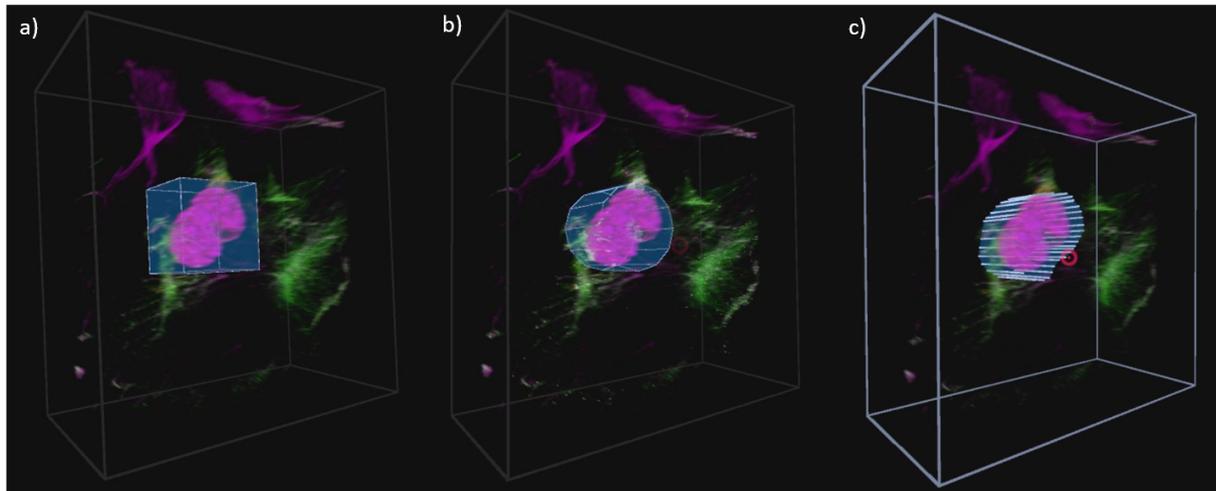


Figure 3.4: The three implemented ROI selection tools. a) The box selection tool, b) the cylinder selection tool and c) the freehand selection tool. Each tool has been used to select a similar part of the volume sample.

When the box and cylinder tools are selected, they initially include the entire volume sample. They must subsequently be transformed to include only the section desired. When the freehand tool is selected, the user traces out the ROI using either head movement (with gamepad) or by pointing the index finger (with hand tracking). Since inaccurate freehand selections may result from the parallax effect, the volume sample is flattened and rotated to face the user before selection. Once the initial rough ROI has been selected, the user has the option to scale the ROI selection independently along any of the three axes. In the case of the freehand tool, the volume is first un-flattened. By scaling along the z-axis, the user can accurately position the ROI within the volume.

After the region of interest has been selected, a ROI mask is generated, in the form of a two-dimensional boolean array, as well as back and front z-positions that indicates at which depths the ROI selection starts and ends.

3.2.5 User trials

Since the average user is generally unfamiliar with VR, we wanted to establish the degree of ease with which users could use our interface. More particularly we wanted to establish how difficult it is to perform certain defined manipulations using the different VR interfaces. Accordingly, user trials were carried out with a diverse group of 29 computer users. Of these, 5 were biologists that regularly work with biological visualisation tools, 15 were engineers (mostly electrical and electronic engineers), another 6 were people with tertiary education in other fields and 3 had no tertiary education. Subjects were between the ages of 21 and 60. Of the participants, 20 were male and 9 female. Many statistical tests are based on the assumption that the data is normally distributed. For this reason we computed the statistical power of the task results using Matlab, to determine whether the rejection of the null hypothesis H_0 is valid. Furthermore, the power analysis proved

that our sample size was adequate to ensure a statistical power of greater than 0.8 for all the tasks, except for the task to transform the sample (which was almost identical for the two interfaces) and the task to scale the ROI along the z-axis (which requires a sample size of 37).

Each participant in the study was asked to perform the same tasks, in the same order, using both the gamepad and the hand tracking interfaces. The time taken to perform each task was measured to allow subsequent objective productivity comparison. Lastly, the participants were asked to perform similar interactions with a traditional keyboard and mouse, without their time being taken, in order for the participants to gain understanding in the current standard interfaces. After using each interface, the participants were asked to complete a subjective questionnaire describing their experience when performing the different tasks.

In order to ensure that the results from different participants were comparable, all the tasks were supervised by the same researcher and all participants followed the same procedure for corresponding samples. The order in which the interfaces were tested was the same for all participants using the gamepad first, then the hand tracking and finally the traditional input. The supervising researcher ensured that all sample transformations, GUI interactions and ROI selections were completed with the same accuracy. The selection accuracies that was required for acceptance is illustrated in Figure 3.4. Furthermore, the users were required to make the selection with the same accuracy between interfaces.

Pre-test preparation

Tests were carried out in a quiet studio environment with only the participant and the researcher present. Each user was informed about the purpose of the test and what they will see and experience. Each participant was then asked to give a subjective rating between 1 and 5 indicating their general computer proficiency, and their experience with biological visualisation and colocalisation analysis. Each participant also provided self-assessments of their experience in using a gamepad, a hand tracking device and a VR headset.

Overall the participants indicated that they had medium to high computer proficiency (Figure 3.5a). The gamepad experience among the participants was diverse (Figure 3.5b) and was a factor that we considered in our later analysis. Most of the participants had very little or no prior exposure to either VR or hand tracking.

Because most participants had no prior VR or hand tracking interface experience, we used two demonstration programs to familiarise them with movement and hand interaction in a VR environment. A stationary VR demo environment and the Blocks demo [47] created by Leap Motion were used for this purpose. Most users are astonished when using VR for the first time, and these introductions helped to ensure that the subjective feedback was based on the effectiveness and productivity of the implementation rather than the initial enthusiasm provoked by VR.

After the participant was comfortable in the VR environment, they were given a brief demonstration of the tasks that they were expected to perform, using our software, as well as an explanation of how to use the given interface. They were then given approximately 10 minutes to familiarise themselves with the interface. Only after they felt comfortable with the interface were they asked to perform the defined tasks, which are described in the next section.

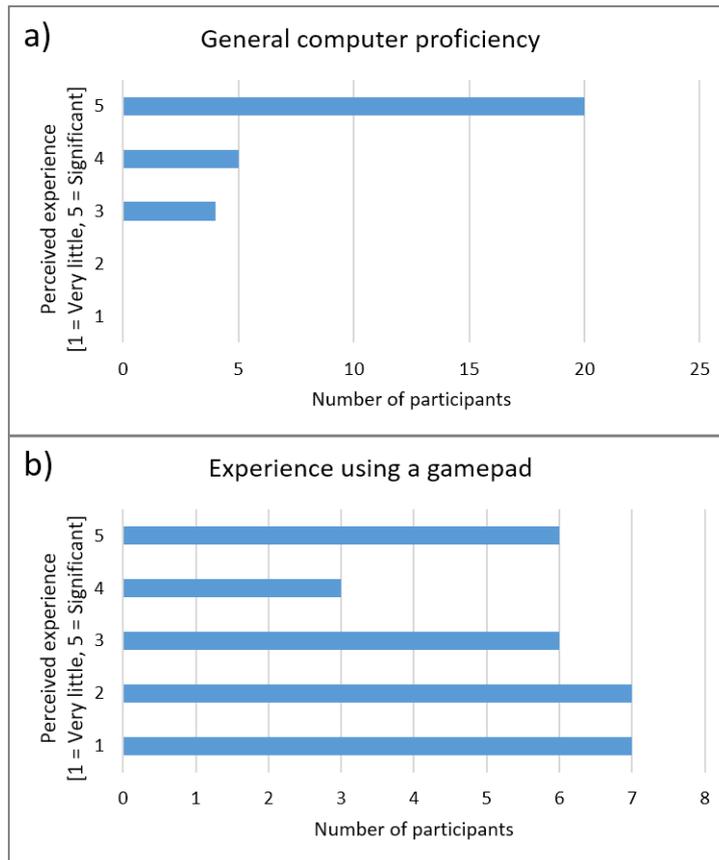


Figure 3.5: Self-assessment of computer proficiency and prior gamepad experience by the subject. a) Participants' general computer experience, and b) gamepad experience.

Objective evaluation

In order to ensure a fair comparison between the participants' experience of and productivity with each VR interface, a series of tasks was devised to ensure that the different aspects of their interactions could be tested and timed. Tasks were chosen that would cover all the aspects of the system that a biological investigator would use to perform a basic sample visualisation and colocalisation analysis:

1. The participants were shown an image of a desired transformation on a volume sample. In order to match this image, the participant needed to perform translation, scaling and rotation of the sample.
2. The participants were asked to change several rendering parameters to prescribed values. This mostly involved changing slider values in the GUI.
3. The participants were asked to place a ROI selection box around a prominent colocalised feature in the sample. This was divided into two steps:
 - a) The ROI box needed to be scaled and translated to the correct position to surround the feature.
 - b) Subsequently the sample needed to be rotated by 90° and the box scaled along the z-axis to match the depth of the colocalised feature.
4. Finally the participants were asked to use the freehand ROI selection tool to accurately outline the colocalised feature and adjust the z-dimension of the ROI.

Each task was explained verbally to each participant immediately before it was performed. The actions of each participant were recorded using both a video camera as well as screen capture. This allowed the time taken for each task to be accurately and unobtrusively measured later. This approach proved to be very effective in making the users feel relaxed while performing the tasks. The entire procedure took between 30 and 45 minutes for each participant. Times were measured to the closest second.

Subjective evaluation

After completing each task, the participants gave a perceived ease of use rating for each interface using a 5 point scale. They were also asked to provide subjective ratings describing how often they forgot the button allocations or the required input gestures. Lastly, they were asked to rate their general sense of how productive they would be when using the interface for microscopy data visualisation and colocalisation analysis.

Since our system was conceptualised with VR at its core, no keyboard or mouse support was implemented. In order to gain insight into how the users perceived the VR interface when compared to conventional input methods, they were asked to perform similar interactions using traditional 3D software using a keyboard and mouse, without the VR headset. The interactions were designed to mimic those implemented in standard biological applications, such as ZEN image software by ZEISS. Subsequently they were again asked to rate the ease of use.

Finally, the participants ranked the three interfaces according to their preference for performing colocalisation analysis, and to how difficult the three interfaces were to learn to use.

3.3 Results and Discussion

Using main effects plots, which compare the means of parts of the data based on different factors, it was found that prior experience of biological visualisation tools or of performing colocalisation analysis had no significant influence on performance. Furthermore, the participants' level of computer proficiency also did not affect their performance. The participants' prior gamepad experience did however influence their performance positively in the gamepad tasks. The average total time taken to complete the required manipulations was more than 12 seconds less for users with significant gamepad experience compared to those with very little experience. The main effects plots of the total time for each interface are shown in Figure 3.6.

3.3.1 Subjective test results

The subjective test results were statistically analysed using contingency tables with Pearson's chi-square (χ^2) test to determine whether differences in participant feedback were statistically significantly ($\alpha = 0.05$). A contingency table is a method of analysing data in a two-way classification table and is usually used when both the dependent and the independent variables are discrete attribute data. It is a tool that can be used to test the relationship between two sources of variation [48]. Contingency tables are therefore useful when analysing our subjective questionnaire feedback, where the independent variable is the interaction method - either the gamepad, hand tracking or keyboard and mouse, and the dependent variable is the participants' subjective rating between 1 and 5. However,

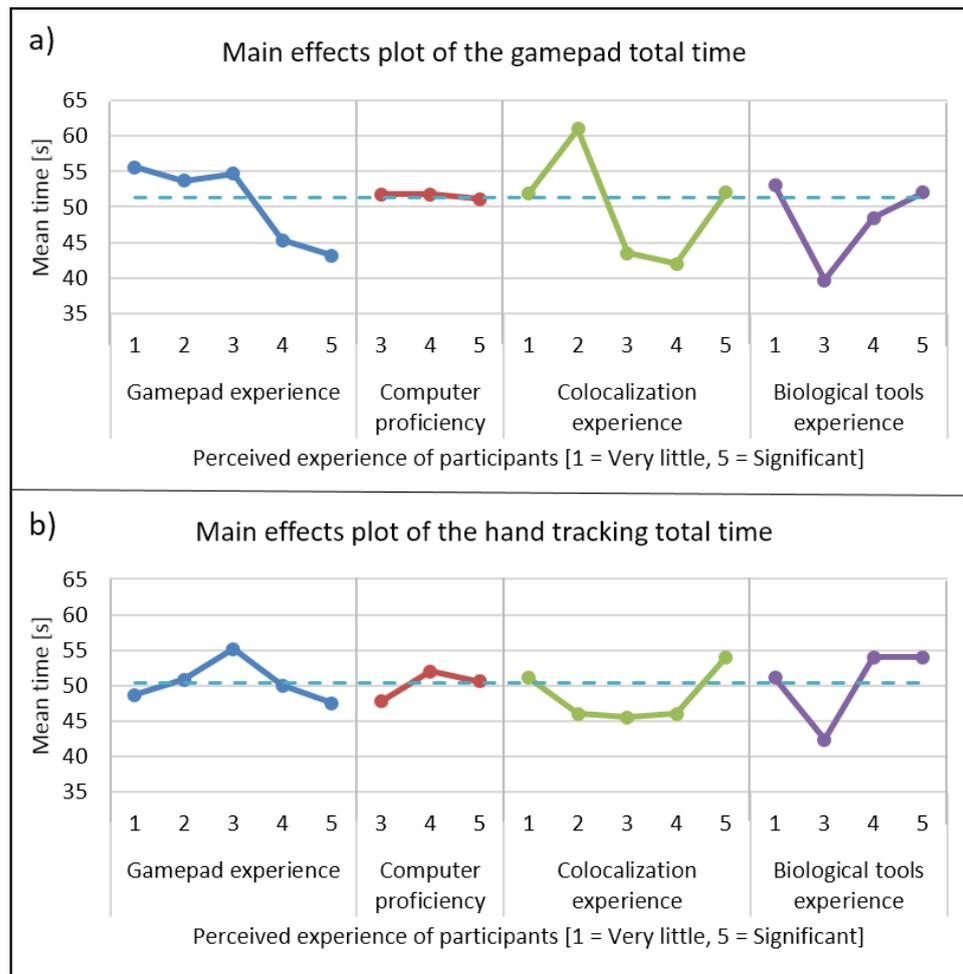


Figure 3.6: The main effects plots for time taken when using gamepad and hand tracking, for a) the gamepad total time, and b) the hand tracking total time for the four categories of self-assessed experience. The dotted line across the plot indicates the overall average time.

since feedback values of 1 and 2 were almost never assigned, these were grouped with the feedback value of 3, in order to allow legitimate application of Pearson's χ^2 test [49].

By taking the average rating given for each of the tasks, it is possible to determine which interface was the preferred method of performing a certain interaction. However, these averages do not reflect whether the participants' perception of two of the interfaces were independent. To determine whether the ratings assigned to each interface differed to a statistically significant degree, we used a χ^2 probability (p-value) of less than or equal to an α of 0.05. Furthermore, we quantified the degree of dependence using Cramér's V , defined by:

$$V = \sqrt{\frac{\chi^2}{N(k-1)}} \quad (3.6)$$

In Equation 3.6, k is the lesser of either the rows or the columns in the contingency table and N is the sample size [50]. The value of V varies between 0, indicating no association between the dependent and independent variables, and 1, for complete association between the variables. If there is high association, then the two interfaces are statistically different. Since our data has 2 degrees of freedom (v), we applied Yates' correction for continuity to the χ^2 calculations before calculating V . In our case a V -value of approximately 0.4

corresponded to a p-value of 0.05. Therefore, V values greater than 0.4 are an indication that the two interfaces are statistically different.

The results of the contingency table analysis are shown in Table 3.1. The average rating for the interface is shown on the diagonal, while p-values are shown below the diagonal and the related V -value are above the diagonal.

Discussion of subjective results

Many of the results for the hand tracking were negatively influenced by inaccuracies in the hand tracking process which are due to current technological limitations, as discussed earlier. Consequently, the user experience can be expected to improve as the technology develops. This should be borne in mind when interpreting the results.

From the results in Table 3.1, for all three transformation interactions (translation, scaling and rotation) the participants gave the highest average rating to the gamepad and the lowest average rating to the traditional interface. Based on the p-values, the participants' experience of the gamepad was statistically different to that of the hand tracking (all p-values < 0.011) and the traditional input (all p-values < 0.001). However, no statistically significant difference between the experience of the hand tracking and traditional input could be established (all p-values > 0.17).

In terms of fitting an ROI box around the colocalised area, no statistically significant preference among the different interfaces could be established (all p-values > 0.12). However, the ratings for scaling the box in the z-axis was significantly different for all three interfaces (all p-values < 0.02), with the gamepad and hand tracking being most similar ($V = 0.894$). In both cases the participants preferred using the gamepad, and most disliked the traditional input.

Many participants found it challenging to use the freehand ROI selection tool using the gamepad with the head tracking. On the other hand, many participants commented that this tool was significantly easier to use with the hand tracking. This is also reflected in the average rating (gamepad $\bar{x} = 3.72$; hand tracking $\bar{x} = 4.66$), where the hand tracking was the preferred interaction method, with the gamepad being the second preferred interface. From the V -values, it is also clear that the hand tracking was experienced significantly different to both the gamepad ($V = 0.632$) and the traditional input ($V = 0.811$). Most users were very frustrated with the inaccuracy of using a mouse to draw the ROI.

Our GUI is based on traditional sliders and buttons, which are familiar from mouse interactions. When the participants were asked to change rendering parameters using the GUI, the gamepad received a higher average rating than the traditional input (gamepad $\bar{x} = 4.83$; traditional input $\bar{x} = 4.69$). However, this difference was not statistically significant ($p = 0.378$). The perception of the hand tracking was statistically significantly worse than the gamepad ($p < 0.001$) and traditional input ($p = 0.005$), with an average of 4.10. Participants felt that interacting with GUI elements using the hand tracking was cumbersome, mainly because it required very precise interactions which were challenging due to inaccuracies in the hand tracking. Furthermore, interacting with GUI elements required the user to hold their hand in the air for long periods of time. This is because the GUI remains stationary relative to the origin in the virtual environment, if the user moved a bit away from the GUI, they would either have to interact with their arms outstretched or first move closer to the GUI. Many users reported that their arms became fatigued after about 10 minutes of use. This indicates that further work in developing GUI interactions that are more suitable for VR would be beneficial. In order to allow direct comparisons, we designed our GUI to be accessible using with both gesture and

Table 3.1: Contingency table summary of subjective results. The average rating per question is given on the diagonal. The highest average is highlighted in green, the second highest average is highlighted in orange and the lowest average is highlighted in red. The p-values are shown below the diagonal and the V-values are shown above the diagonal. The p and V-values that are statistically significant are highlighted.

Sample Transformation			
Translation	GP	HT	KM
Gamepad (GP)	4,83	0,472	0,747
Hand Tracking (HT)	0,011	4,17	0,282
Keyboard and Mouse (KM)	<0,001	0,171	3,52
Scaling	GP	HT	KM
Gamepad (GP)	4,83	0,642	0,665
Hand Tracking (HT)	0,001	3,93	0,075
Keyboard and Mouse (KM)	<0,001	0,792	3,79
Rotation	GP	HT	KM
Gamepad (GP)	4,66	0,605	0,796
Hand Tracking (HT)	0,002	3,72	0,171
Keyboard and Mouse (KM)	<0,001	0,415	3,17
ROI Selection			
Moving ROI Box	GP	HT	KM
Gamepad (GP)	4,59	0,207	0,295
Hand Tracking (HT)	0,313	4,28	0,116
Keyboard and Mouse (KM)	0,126	0,603	4,21
Scaling ROI along z-axis	GP	HT	KM
Gamepad (GP)	4,69	0,442	0,894
Hand Tracking (HT)	0,018	4,10	0,571
Keyboard and Mouse (KM)	<0,001	0,003	3,03
Freehand ROI tool	GP	HT	KM
Gamepad (GP)	3,72	0,632	0,306
Hand Tracking (HT)	0,001	4,66	0,811
Keyboard and Mouse (KM)	0,143	<0,001	3,31
General experience			
Change GUI settings	GP	HT	KM
Gamepad (GP)	4,83	0,714	0,235
Hand Tracking (HT)	<0,001	4,10	0,526
Keyboard and Mouse (KM)	0,378	0,005	4,69
Required input doubt	GP	HT	KM
Gamepad (GP)	3,62	0,860	0,261
Hand Tracking (HT)	<0,001	4,83	0,851
Keyboard and Mouse (KM)	0,223	<0,001	3,52
General productivity	GP	HT	KM
Gamepad (GP)	4,66	0,523	0,718
Hand Tracking (HT)	0,007	4,03	0,178
Keyboard and Mouse (KM)	<0,001	0,415	3,59

gamepad interfaces. A gesture-only system need not be constrained in this way. One common alternative implementation for simple GUI interfaces is to attach the GUI to one hand and interact with it using the other hand. Our GUI was too complex for this design to be practical, however.

When the participants were asked how frequently they forgot which interaction (button press or hand gesture) performed which function, as might be expected, most felt that they never forgot the required interaction for the hand tracking. In fact it was often commented that the hand tracking was very intuitive to use and required almost no explanation. However, since button allocations had to be memorised for both the gamepad and traditional input, the initial cognitive load was greater and users had a similar experience with both. It was, however, generally felt that if either the gamepad or traditional interfaces were used for a few hours, the button allocations would be remembered instinctively. Therefore, in the long term, the intuitiveness of the hand tracking may not remain a significant advantaged.

Lastly, when the participants were asked how productive they thought they would be when performing microscopy data visualisation and colocalisation analysis with each interface, the gamepad received the highest average rating (4.66) with the hand tracking receiving the second highest (4.03). The main reason furnished for this response was the perceived inaccuracies in the hand tracking (described earlier), which made it difficult to use for some users.

Subjective interface rankings

When the participants were asked which interface they would prefer for colocalisation analysis, 15 indicated the gamepad, 11 indicated hand tracking and the remaining 3 indicated traditional input, of which 2 were biologists, mainly because of its familiarity. A summary of the rankings is shown in Figure 3.7a. A recurring reason that was given to explain why the hand tracking was not the preferred choice was that the hand tracking system was too inaccurate for fine interactions, which made certain tasks more challenging than they should be. A general feeling was that if the hand tracking was improved, it might be preferred.

The participants were also asked about their perception of how easy each interface would be to learn to use. 21 of the participants felt that the hand tracking was the easiest to learn, since it was the most intuitive to use. 7 participants experienced the gamepad interface as the easiest to learn and one person preferred the keyboard and mouse, due to its familiarity. This does however show that the gamepad is perceived to be easier to learn than an already familiar interface. A summary of these rankings is shown in Figure 3.7b.

3.3.2 Objective test results

The time the participants took to perform each task was statistically analysed using the t-test. The t-test is used to determine whether two sets of data are significantly different from each other. The results of its application are shown in Table 3.2. Comparative box plots for the two interfaces for all the tasks are shown in Figure 3.8. The t-test is based on the assumption that the data are normally distributed. Using the Anderson-Darling normality test it was determined that all gamepad tests were normally distributed ($p > 0.05$). Interestingly, except for the tasks in which the parameters were changed using the GUI and in which the ROI box was scaled along the z-axis, none of the hand tracking timing results were normally distributed ($p < 0.05$). However, the total time for all the

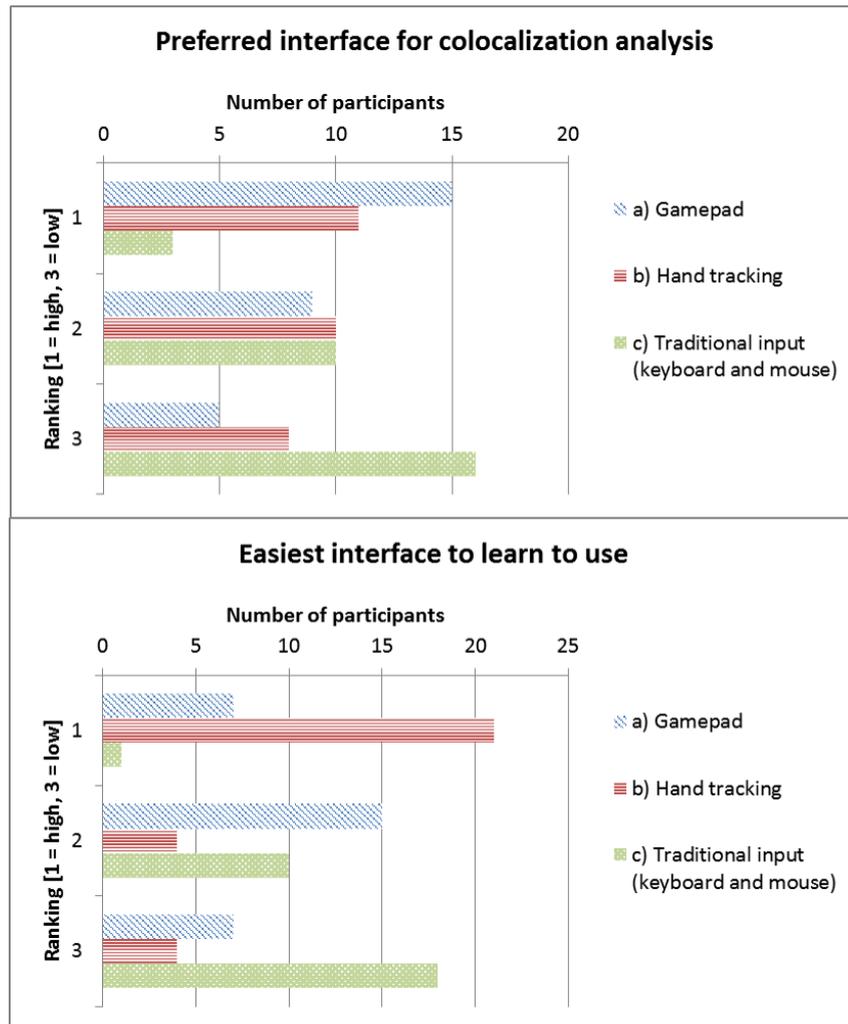


Figure 3.7: Subjective ranking of interfaces. a) The preferred interface to use for colocalisation analysis, and b) the easiest interface to learn to use.

tasks using the hand tracking was again normally distributed. The deviation from the normal distribution of some hand tracking tasks can largely be explained by delays caused by inaccuracies in the hand tracking itself. This means that the participants did not all share the same experience. Therefore, the t-test results that compare the gamepad and hand tracking should be regarded with some caution, even though the t-test is known to be robust also for non-normal data [51].

Discussion of objective results

For all tasks, the average time taken to perform the required action was shorter for the hand tracking than for the gamepad, by a cumulative average of 3 seconds. Only the task in which the participants were asked to change the rendering parameters using the GUI did the gamepad perform better by an average of 2.9 seconds. As stated earlier, the inaccuracies in the hand tracking as well as GUI elements that are not optimised for hand interactions led many users to experience gesture interaction with the GUI as cumbersome.

Since the participants used the interfaces in the same order, a learning effect, in which participants became more familiar with the visualisation tools and tasks, is in

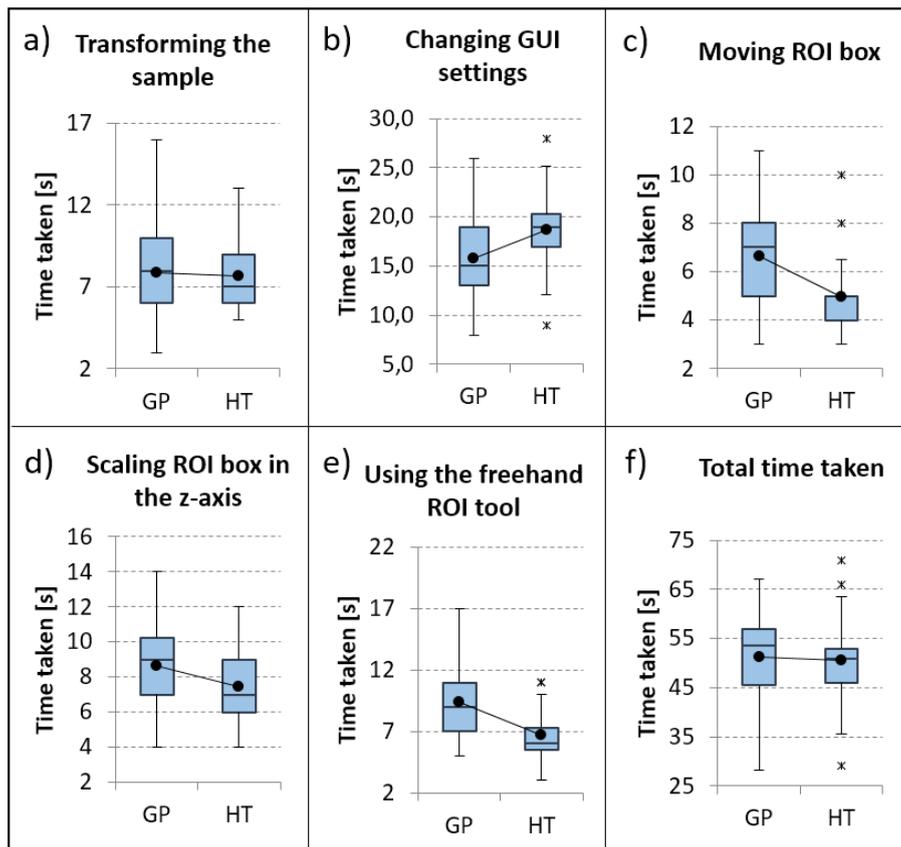


Figure 3.8: Comparative box plots for objective tests. Box plots for the objective tests that comparing the gamepad (GP) and hand tracking (HT) interfaces.

principle possible. This could, in turn, have contributed to the improved performance of the hand tracking. However, the hand tracking and gamepad interfaces differ vastly in how interactions are performed and each participant were thoroughly familiarised with the hardware and software environment before the test. Hence we do not believe that the average times we measured are significantly influenced by a learning effect.

Based on the p-value associated with the t-test, the gamepad provided a benefit when changing the GUI settings ($p < 0.01$). When moving the ROI box ($p < 0.01$) and using the freehand tool ($p < 0.01$) the hand tracking performed statistically significantly better than the gamepad. Even though the task of scaling the ROI along the z-axis indicates differences between the interfaces ($p = 0.03$), this was inconclusive since there is a Beta-risk of 30%. There is also no clear benefit to using either interface for the task of transforming the sample ($p = 0.38$). This is also suggested by the similar average times (gamepad $\bar{x} = 7.9$; hand tracking $\bar{x} = 7.7$), as well as the low statistical power of 0.0655.

From Table 3.2, the standard deviation of all timing results was smaller for the hand tracking than for the gamepad, except for the GUI interaction task. This can largely be ascribed to the fact that the hand tracking was more intuitive to use and was quickly learnt. On the other hand, the large standard deviation for the gamepad indicates that some participants were able to use the gamepad more productively with less training than others. This may be a result of prior gamepad experience.

Table 3.2: Comparison of objective results. The average times (in seconds) and standard deviation for all the objective tasks for both interfaces are shown, as well as the statistical power values for each task. The p-value from the t-test indicates a statistically significant difference between interfaces when $p < 0.05$. All statistically significant differences are highlighted.

Task	Gamepad		Hand tracking		t-test (p-val)	Power (1- β)
	Mean	σ	Mean	σ		
Transforming the sample	7,9	2,8	7,7	2,0	0,38	0,0655
Changing GUI settings	15,8	3,7	18,7	3,9	<0,01	0,9966
Moving ROI box	6,7	1,9	5,0	1,4	<0,01	0,9992
Scaling ROI along z-axis	8,7	2,5	7,4	2,1	0,03	0,6973
Freehand ROI tool	9,4	2,7	6,7	2,1	<0,01	0,9800
Total time taken	51,3	10,2	50,4	8,1	0,36	0,0723

Influence of prior gamepad experience

In terms of prior gamepad experience, the participants can be divided into two groups - those who are experienced in using a gamepad and those who are not. There were 14 participants that gave a low rating (1 or 2) for gamepad experience and 15 participants that gave a high rating (3, 4 or 5). In order to determine to what degree gamepad experience influenced performance, a brief comparison was made between these two groups.

From the analysis it became clear that, on average, participants with high gamepad experience performed better for every task when using the gamepad than those with low gamepad experience. However, except for the task in which the ROI box was scaled along the z-axis, none of these differences are statistically significant (all p-values > 0.065 except for scaling the ROI along the z-axis $p = 0.005$). Comparative interval plots for gamepad usage are shown in Figure 3.9. From these plots it is also clear that high gamepad experience did not consistently decrease the variability in the data. Furthermore, the overall standard deviation is significantly higher ($\sigma = 11.85$) for users with high gamepad experience than those with low gamepad experience ($\sigma = 7.07$). This indicates that not all users received the same productivity benefit from familiarity with this input devices. In general, there seems to be no conclusive influence of gamepad experience on hand tracking performance.

3.3.3 Contribution of system to improved analysis

The user perception in data handling is crucial to sample control, and can allow subsequent enhanced analysis selectivity. In colocalisation analysis, it is particularly true for maximising control over ROI selection and the associated thresholds. Often such data analysis is not performed properly due to challenges in working with the available software tools. Our system aims to allow better visualisation and manipulation of colocalisation data in order to improve such analysis.

The user study indicates that our system successfully offers ease of use of powerful analysis tools. This has the potential to improve the quality of colocalisation analysis.

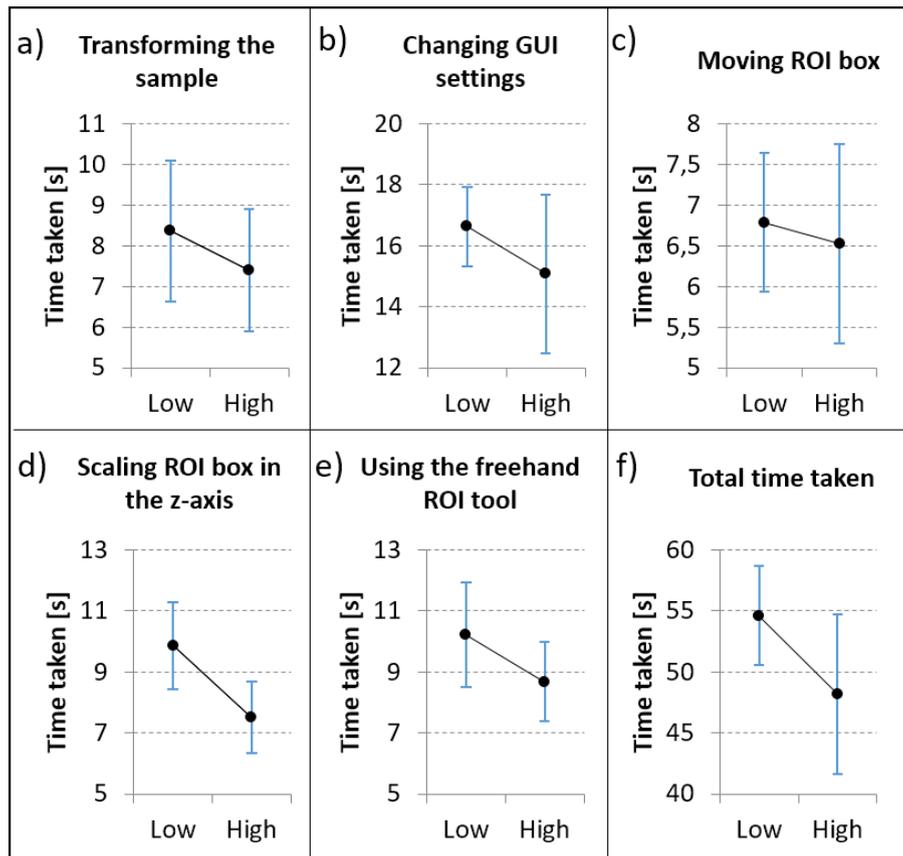


Figure 3.9: Comparative interval plots for low and high gamepad experience, showing the mean and 95% confidence intervals. Participants that gave a gamepad experience rating of 1 or 2 were considered to have low experience, while a rating of 3, 4 or 5 were considered as high experience.

3.4 Conclusions

We have presented a system for visualising three-dimensional microscopy data in a virtual reality (VR) environment. As a specific objective we have focused on the calculation and analysis of colocalisation within this VR environment. Our system offers fully-immersive manipulation of the microscopy data and analysis tools by means of an intuitive hand gesture interface. As an alternative method of input, a conventional gamepad was also employed. A set of user trials was performed to determine the effectiveness of our system when performing colocalisation analysis and related manipulations in the VR environment (see supplementary video at <https://www.youtube.com/watch?v=ajdNKnAHFMw>).

Overall, users were very enthusiastic about the practical possibilities of the system. Interviews conducted with the five biological experts among these users, indicated a strong conviction that this system would greatly benefit them in their current work in terms of insight into the sample, as well as analysis productivity. Objective measurements show that, for most tasks, users were most productive, to a statistically significant degree, using the hand tracking interface. Despite this, the subjective assessments indicated that most users were under the impression that they were more productive using the gamepad. This negative subjective bias towards the hand tracking is ascribed to current inaccuracies in the hand tracking process, which often led to frustration. It is likely that, as hand tracking technology improves, this bias will diminish and hand tracking could become the preferred option for interaction in the VR environment. Indeed, the user trials showed that hand

tracking is already perceived as the most intuitive method of interaction.

A secondary conclusion is that more research is required into the design of effective graphical user interface (GUI) for hand gesture input in a VR environment. The familiar offering of buttons and sliders was found by many users to be cumbersome, difficult to use with precision, and sometimes leading to fatigue.

In the next chapter we apply this system to the analysis of new biological samples, specifically the precise 3D ROI selection with subsequent 3D colocalisation analysis. We compare our results with the traditional 2D analysis that is performed using maximum intensity projections.

Chapter 4

Improved ROI selection and colocalisation analysis in 3D fluorescence microscopy samples using virtual reality

The content of this chapter was published as the following research article:

Theart, R.P., Loos, B., Powrie, Y.S. and Niesler, T.R., 2018. Improved region of interest selection and colocalization analysis in three-dimensional fluorescence microscopy samples using virtual reality. *PLOS one*, 13(8), p.e0201965.

In the previous chapter, we demonstrated how our virtual reality based system can be applied to the analysis of colocalisation in fluorescence microscopy z-stacks. In this chapter we will investigate this further by determining the benefits that performing region of interest selection and subsequent colocalisation analysis in three-dimensions offers over the traditionally used two-dimensional maximum intensity projection analysis.

4.1 Introduction

Fluorescence microscopy allows the tagging of particular macromolecules of interest and the subsequent study of their role in biological processes [52, 53]. With advances in fluorescence image acquisition and processing, such as confocal and Super Resolution Structured Illumination Microscopy (SR-SIM), it has become common practice to analyse the z-stacks, consisting of multiple two-dimensional (2D) fluorescence image frames, which these microscopes produce. These 2D images can be three-dimensionally (3D) reconstructed using direct-volume rendering [20, 22, 54, 55], allowing the interrogation of the sample in 3D which can provide additional insight into biological events.

In the biological sciences, the interaction of two molecular structures or proteins is often of key interest, and hence colocalisation analysis is commonly performed, as was discussed in Sections 3.1 and 3.1.1. In summary, colocalisation refers to the geometric codistribution of two fluorescent labels or colour channels. It is usually assessed over a whole cell and metrically described using Pearson's correlation coefficient (PCC), Manders' overlap coefficient (MOC) and Manders' correlation coefficient (MCC). In order to perform accurate colocalisation analysis, it is important to precisely identify and select

the region of interest (ROI) over which the colocalisation will be calculated. The ROI is a sub-volume of the sample that is considered in the analysis. However, the precise selection of the ROI within a sample of interest, in 3D space, remains a practical challenge. Although the z-stack is inherently 3D, most established colocalisation analysis methods define the ROI on a 2D projection of the z-stack, usually the maximum intensity projection (MIP). The MIP is a data visualisation method in which the voxel with the maximum intensity along the view direction is projected onto a 2D plane. In the case of a z-stack, the view direction is usually aligned with the z-axis. Because the projection is usually orthographic, at each x-y coordinate the MIP retains only the voxel among all the image frames with the maximum intensity. Using this projection, however, excludes most of the data from the analysis. Furthermore, complex 3D structures cannot, in general, be adequately selected on the 2D projected image. Consequently, the use of the MIP for ROI selection can result in both the exclusion of relevant data points and the inclusion of irrelevant data, leading to possible inaccuracies in the colocalisation analysis.

Often, the precise analysis of sub-cellular structures is hampered by the limitations of current ROI selection methods. It has been shown that a 3D mask can be used to make a 3D selection [56, 57, 58]. The most common way this is done is by drawing boundary lines around the sample to create a mesh. This mesh is then allowed to shrink until it fits closely around the structure requiring analysis, after which every data point inside the mesh forms part of the ROI [59]. This approach works well for clearly defined regions, but it is not interactive and does not easily extend to complex selection combinations. Moreover, when using conventional 3D reconstruction and visualisation of a sample, this reconstruction is ultimately displayed on a 2D screen which often leads to ambiguous data interpretation, especially when the 3D structure in the sample is complex. An inaccurate interpretation of the visualisation can lead to inaccuracies in the selection, since it is not immediately apparent how well the 3D selection matches the desired ROI, without constantly rotating the sample, or providing several view angles.

We employ the virtual reality (VR) system that was developed in Chapters 2 and 3 which allows sample visualisation and ROI selection in true 3D, thereby achieving an unambiguous representation of the data. Since more data points are taken into account when using a 3D ROI, the statistical uncertainty associated with the colocalisation metrics will also be reduced [60, 61]. We use this VR-assisted 3D ROI selection method to analyse a neuronal cell that is obtained using SR-SIM. By selecting regions and sub-cellular structures of biological relevance, such as the perinuclear, middle and membrane regions and performing extensive colocalisation analysis, we demonstrate the improved precision and control achieved by the 3D ROI selection. Neuronal cells have been chosen as the subject of our study due to the particularly high relevance of accurate ROI control to their analysis. By assessing distinct structures, such as the tubulin network and protein aggregates, we demonstrate how the 3D ROI selection can enhance our understanding of biological processes. Specifically, we have assessed the dissociation between the protein Tau and the microtubulin network in an *in vitro* model of autophagy dysfunction, which mimics a key pathology in Alzheimer's disease [62], by analysing the colocalisation of Tau and acetylated α -tubulin.

4.1.1 Neuronal cell investigation

One of the key structures of interest in neurobiology is the microtubule network, not only due to its critical role in neuronal integrity and transport processes, but also due to its role in the development of neuronal pathology. Alzheimer's disease is a neurode-

generative disorder of the brain and a leading cause of dementia. Severe cognitive and short term memory deficits are commonly associated with this disease. The pathology is characterised by protein aggregates that manifest in brain tissue due to a dysfunction in protein degradation systems. One such affected system is the autophagy lysosomal pathway (ALP), which is a major catabolic process in eukaryotic cells that is responsible for the degradation of long-lived or damaged proteins and organelles. This pathway relies on the microtubulin (MT) network for the transportation of autophagosomes, which are double-membraned structures that enclose the proteins or organelles destined for degradation, towards lysosomes, which are compartments (or vesicular structures) within cells that contain enzymes with acidic properties. The lysosomes fuse with the autophagosome and degrade the constituents of the vesicle, now known as an autophagolysosome. Disruption of the ALP is a known perturbation that occurs within pathogenesis of Alzheimer's disease [63, 64, 65]. The deterioration in the breakdown of autophagolysosomes leads to their accumulation on the MT, which disrupts the normal transport of other vesicles such as mitochondria [66, 67]. The disruption of the ALP further leads to the accumulation of proteins that rely on this degradative system to maintain their turnover. Among these aggregates are the neurofibrillary tangles (NFTs) which are composed of the microtubule-associated protein Tau, that controls the stability of the microtubulin network [68]. MT acts as part of the cytoskeleton and platform for dynein dependent autophagosome transport [69]. The fluorescent labels used in our analysis make the Tau visible as a green signal and the MT (acetylated α -tubulin) visible as a red signal in the SR-SIM z-stack. The relationship between autophagy dysfunction and microtubule stability, however, remains unclear. Therefore, the interaction between Tau and MT as well as their colocalisation is of interest and was chosen as the focus of this study.

4.2 Materials and methods

Neurons were cultured under conditions that simulate what is observed in Alzheimer's disease. For that purpose a particular mechanism, protein clearance, was disrupted, using chloroquine. Since Alzheimer's disease is a progressive disease, two time points were chosen, one that simulates early, and one that simulates late stage Alzheimer's disease. In order to visualise specific features of cell disruption, a fluorescing protein, GFP-tau, was used. In order to visualise a second structure in the same sample, the tubulin network, an antibody binding approach, linked to a second, spectrally different fluorescent dye, was employed. Finally, cells were visualised using specialised optics (SR-SIM) and subsequent image processing was performed, to generate an image that achieves sub diffraction resolution. Samples were subjected to the colocalisation analysis. A more detailed description of this process follows below.

4.2.1 Cell culture and transfections

GT1-7 murine hypothalamic cells were cultured in DMEM supplemented 10% FBS and 1% PenStrep at 37°C and 5% CO₂ atmosphere and treated with chloroquine diphosphate. In order to functionally disrupt autophagy, a key feature in neurodegeneration, GT1-7 cells were exposed to 100 μ M of CQ for 6 and 24 hours respectively. 50 000 cells were trypsinised, counted and washed in PBS. Next, cells were resuspended in 100 μ L sterile Neon[®] Resuspension buffer containing 5 μ g of GFP-Tau (P301L) DNA. The suspen-

sion was aspirated into a gold plated Neon[®] Tip using a Neon[®] Pipette and cells were electroporated at 1350 V for 1 pulse lasting 30 ms.

4.2.2 Immunofluorescence and super resolution structured illumination microscopy (SR-SIM)

Transfected cells were seeded into 6-well dishes containing 5 mm round coverslips and 2 mL fresh media treated with chloroquine after 24 hours, followed by fixation for 10 minutes in a 1:1 ratio of 4% Paraformaldehyde: DMEM at 37°C. Cells were blocked using 3% donkey serum in PBS, followed by probing the required primary antibody, i.e. acetylated α -tubulin (Santa Cruz, 23950) overnight at 4°C. Next, cells were incubated with a secondary antibody, Alexa Fluor 568 (ThermoFisher, A-10042) for 50 minutes. Coverslips were then washed 3x5 minutes with PBS and using Dako[®] fluorescent mounting media. For SR-SIM analysis of GFP-Tau transfected cells, thin (0.1 μm) z-stacks of high-resolution (1024x1024 pixels) image frames were collected by utilising an alpha Plan-Apochromat 60x/1.4 oil immersion DIC M27 ELYRA objective on an ELYRA PS.1 system (Carl Zeiss Microimaging, Germany) equipped with a 488nm laser (100mW) in 5 rotations, 561nm laser (100mW) in 5 rotations and an Andor EM-CCD camera (iXon DU 885). Images were reconstructed using Zeiss Zen Black Software (2012) based on a structured illumination algorithm [70]. Colocalisation analysis of these reconstructed super-resolution images using both the 2D and 3D VR-assisted 3D ROI selection systems.

4.2.3 Method

We performed viability assays to determine the degree of cell death after cells were exposed to chloroquine. Two points in time were chosen in order to mimic a mild and more severe neuronal injury, with significant cell death occurring only after 24 hours (Figure 4.1). This would allow the assessment of the Tau-microtubulin interaction in a time and injury dependent manner and consequently the study of the molecular perturbations occurring prior to and at the onset of cell death (i.e. neurodegeneration). Using SR-SIM, cells were imaged at 0 hours (control, before chloroquine exposure) and again after 6 and 24 hours (treated, after chloroquine exposure). We will henceforth refer to these points in time as $t = 0$, $t = 6$ and $t = 24$ respectively. Four representative cells were selected for analysis at each point in time. Hence our data set consisted of a total of 12 cells. For each of these cells the z-stack consisted of 30 slice images on average.

We hypothesised high levels of colocalisation under healthy control conditions, with increasing dissociation between Tau and microtubulin and an associated loss of colocalisation that reach their highest levels at $t = 24$ after exposure to chloroquine. The colocalisation was analysed in both 2D (using the MIP) and 3D for the whole cell, the perinuclear region, the middle and membrane regions of the cell and a few prominent microtubulin strands and protein aggregates. These regions, shown in Figure 4.2, were chosen in order to determine the contribution of each to the colocalisation profile of the cell as a whole.

4.2.3.1 Colocalisation metrics

Colocalisation was calculated by applying three metrics to the 3D z-stack. These metrics were Pearson's Correlation Coefficient (PCC) [71, 72], Manders' Overlap Coefficient

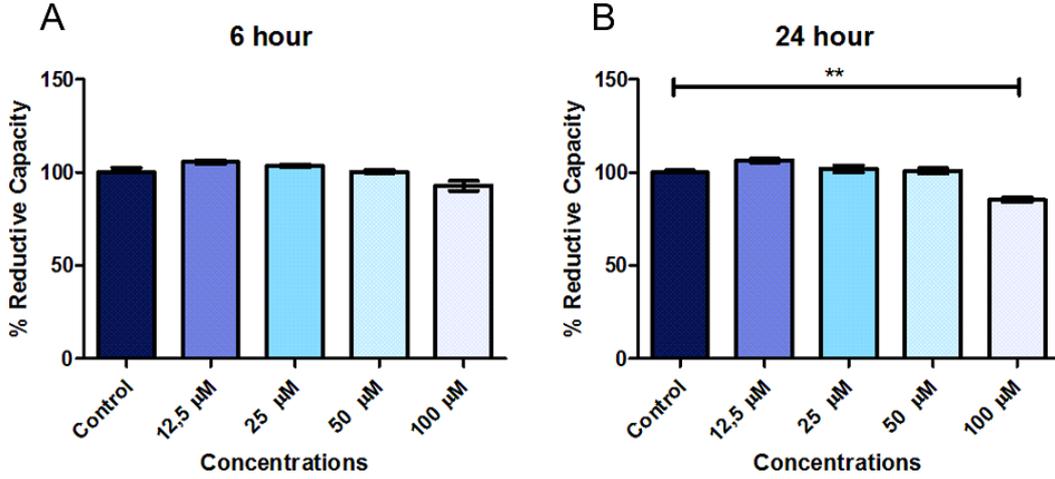


Figure 4.1: WST-1 reductive capacity assay of cells treated with CQ for 6 hours (A) and 24 hours (B), respectively. 100 μM CQ was sufficient to cause a significant reduction in cell viability after 24 hours of CQ exposure, but not after 6 hours. All data are presented as a percentage of control (Mean \pm SEM). ** ($p < 0.05$ vs control). N = 3 independent experiments.

Table 4.1: Colocalisation metrics. R_i represents the intensity of voxels in channel 1 (red fluorescent label) and G_i represents the intensity of voxels in channel 2 (green fluorescent label). \bar{R} and \bar{G} represent the mean intensities of the respective channels [74].

Colocalisation Metric	Equation	Number
Pearson's Correlation Coefficient	$\text{PCC} = \frac{\sum_i (R_i - \bar{R}) \times (G_i - \bar{G})}{\sqrt{\sum_i (R_i - \bar{R})^2 \times \sum_i (G_i - \bar{G})^2}}$	Eq 1
Manders' Overlap Coefficient	$\text{MOC} = \frac{\sum_i (R_i \times G_i)}{\sqrt{\sum_i R_i^2 \times \sum_i G_i^2}}$	Eq 2
Manders' Correlation Coefficient (Ch 1)	$\text{M1} = \frac{\sum_i R_{i,\text{colocal}}}{\sum_i R_i}$ where $R_{i,\text{colocal}} = \begin{cases} R_i, & \text{if } G_i > 0 \\ 0, & \text{if } G_i = 0 \end{cases}$	Eq 3
Manders' Correlation Coefficient (Ch 2)	$\text{M2} = \frac{\sum_i G_{i,\text{colocal}}}{\sum_i G_i}$ where $G_{i,\text{colocal}} = \begin{cases} G_i, & \text{if } R_i > 0 \\ 0, & \text{if } R_i = 0 \end{cases}$	Eq 4

(MOC) and Manders' Correlation Coefficient (MCC) [73, 74]. The equations used to calculate these metrics were presented in Section 3.1.1 and are summarised in Table 4.1.

4.2.3.2 Colocalisation visualisation

It is standard practice to evaluate colocalisation not only by means of metrics, but also by means of visualisation. We therefore also visualised the colocalised volumes by either su-

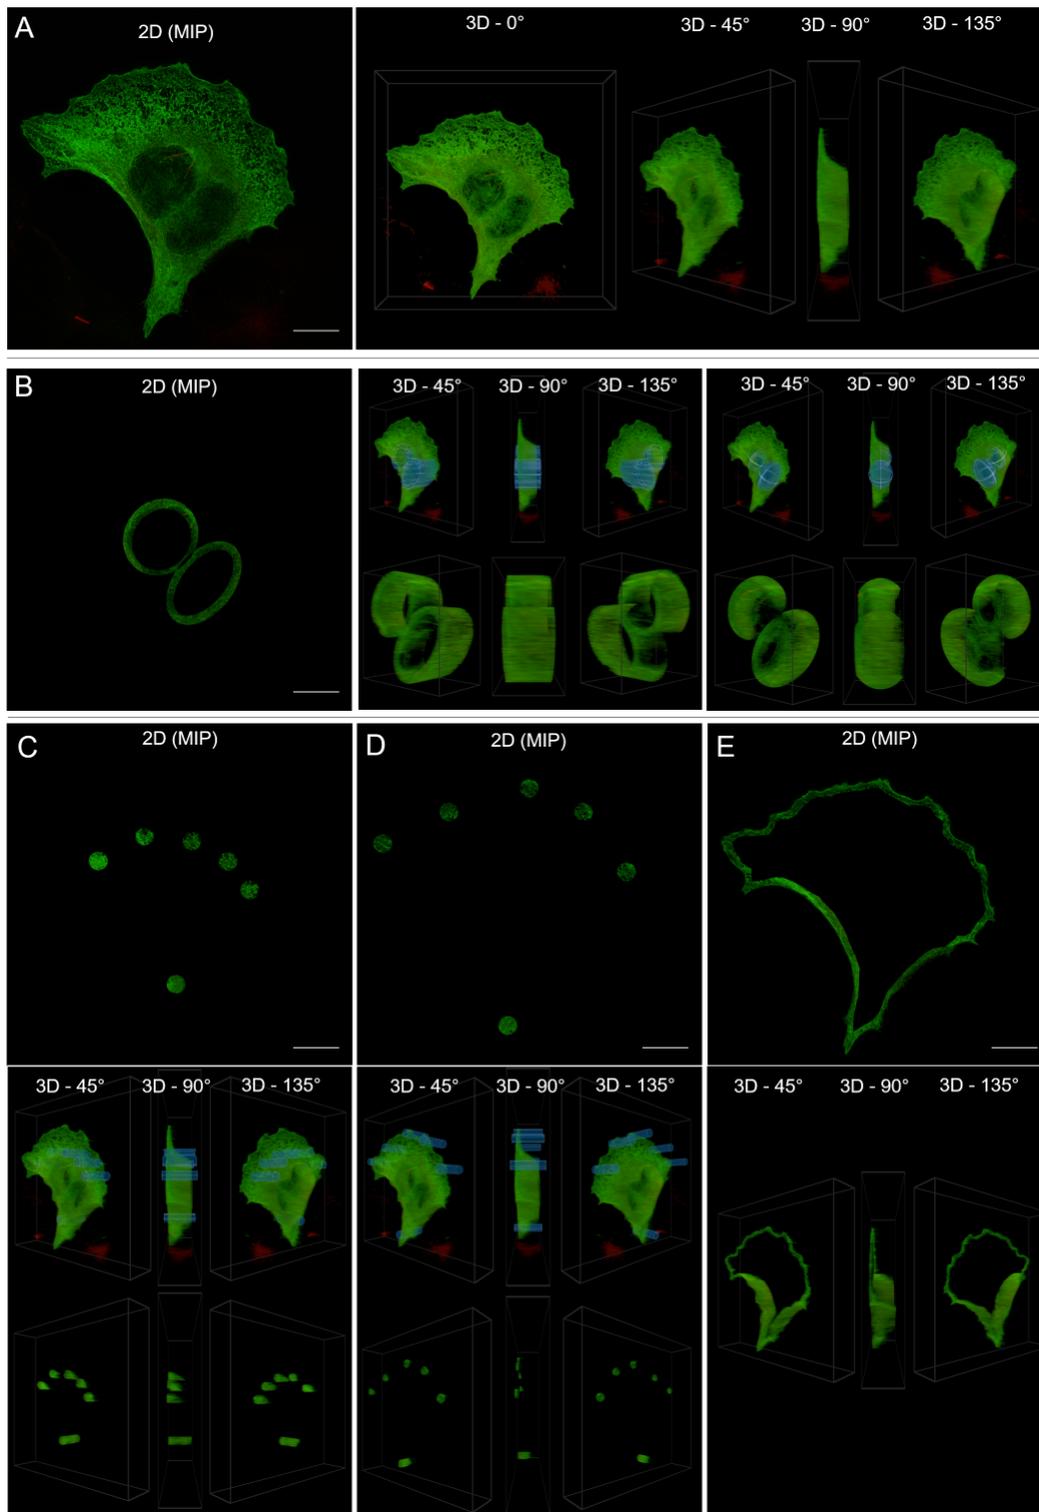


Figure 4.2: Comparison of the MIP and the 3D visualisation of some of the main regions of interest analysed for colocalisation. The 3D image shows the selection overlaid on the whole cell, as well as the selection isolated from the remainder of the cell. The depicted cell is from the control sample set ($t = 0$). A: The whole cell. B: The perinuclear region for both a cylindrical and ellipsoidal selection. C: The sampled middle region. D: The sampled membrane region. E: The freehand membrane region (no overlay of ROI shown).

perimposing them as white voxels on the 3D sample, or by visualising only the colocalised

voxels in white or directly in their fluorescence emission spectrum colours [75, 76]. This served largely to assist in making ROI selections. To allow comparison, the colocalisation was visualised on the MIP in a similar manner.

4.2.3.3 ROI selection tools

Using the VR system, the following 3D selection tools were employed: *box*, *cylinder*, *sphere* and *freehand* (Figure 4.3), or any combination thereof. We will henceforth refer to each application of a tool as a selection. For the freehand selection, ROIs were traced out in 3D using either head movement and a gamepad or alternatively by pointing and drawing with the index finger using a hand tracking system, as has been demonstrated in Chapter 3. It was possible to isolate and enlarge a selection, such as a single tubulin strand, and subsequently refine this selection. In this way smaller, more precisely defined sections of the sample could be visualised and analysed independently. For a demonstration of this process in practice see the supplementary video <https://youtu.be/nC2xh19KLGg>.

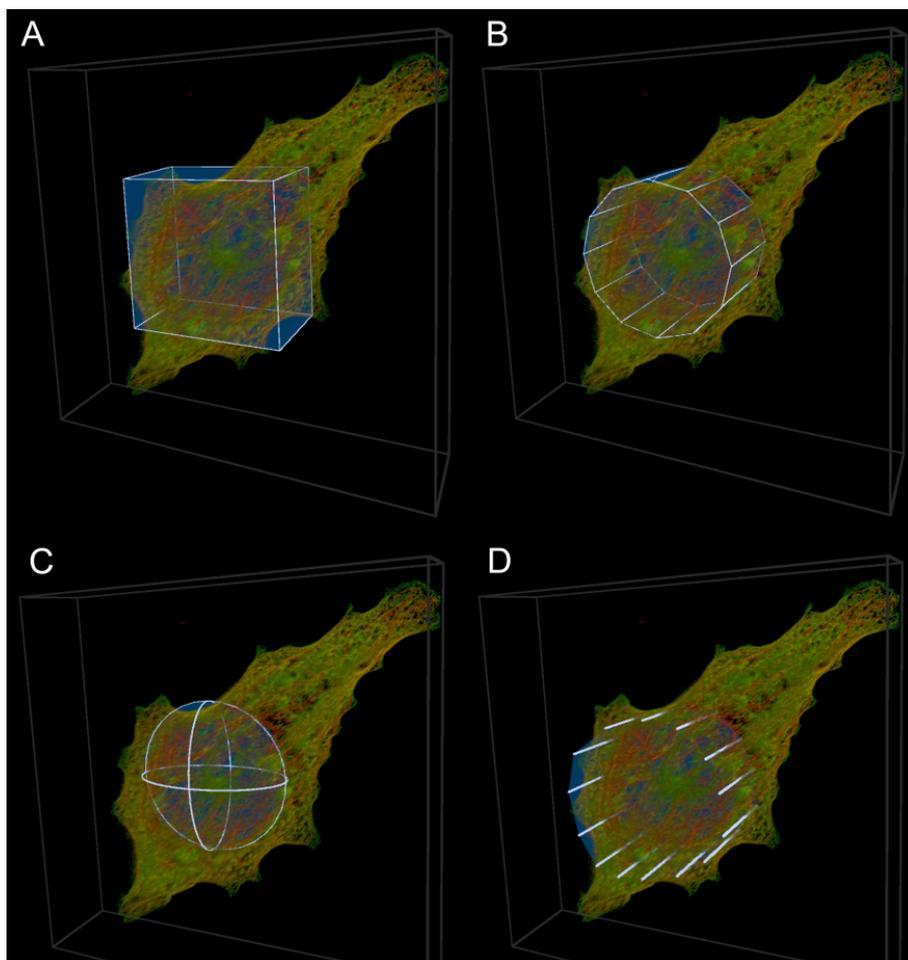


Figure 4.3: The ROI selection tools employed in the analysis. Combinations of these tools were used to make selections. A: The box tool. B: The cylinder tool. C: The sphere tool. D: The freehand tool.

It was also possible to remove the inner part of a selection by using selections within selections, which provided greater versatility. Advantage was taken of this when selecting the perinuclear region of a cell, where the inside of the nucleus needed to be excluded

from the ROI by placing a smaller ellipsoid within a larger ellipsoid and excluding the overlapping volume (Figure 4.2B).

In order to directly compare the 3D and MIP approaches, the closest comparable 2D ROI was used when calculating colocalisation from the MIP. When selecting the perinuclear region, for example, the 3D ROI was an ellipsoid within an ellipsoid, which is most closely represented by an ellipse within an ellipse in 2D. Since an ellipse within an ellipse on the MIP corresponds to an elliptic cylinder within an elliptic cylinder in 3D, we see that certain selections are only possible when specified directly in three dimensions.

4.2.3.4 Analysis

For the colocalisation analysis, the following six ROIs were chosen: the whole cell, the perinuclear region, six to eight evenly spaced ROIs in both the middle and membrane regions of the cell, the entire membrane selected continuously using the freehand tool, microtubulin strands and protein aggregates. In order to gain insight into the effect of the 3D selection geometry on the colocalisation results, both an elliptic cylinder and an ellipsoid were used when making 3D ROI selections. This is especially important for the perinuclear region where the underlying structure resembles an ellipsoid.

A total of four representative cells at each considered time ($t = 0$, $t = 6$ and $t = 24$) were analysed for colocalisation using both the 3D and the 2D MIP approaches to ROI selection. When comparing analysis results between the 3D and 2D selections, we will refer to a comparison of the *ROI selection methods*. For each point in time the mean of each colocalisation metric over the four cells was calculated. A 95% confidence interval was calculated for each mean based on the t-distribution. Since four samples per group were analysed for each estimate of the mean, the associated z value for the 95% confidence interval is 2.776. Note that the confidence interval was not used to determine significant differences but only as an indication of the spread of the data. Confidence intervals were calculated as follows (where s is the sample standard deviation):

$$CI_{95\%} = \bar{x} \pm 2.776 \times \frac{s}{\sqrt{4}} \quad (5)$$

For the sampled (cylinder and ellipsoid) membrane regions at $t = 24$, the sample size was only two and not four, since in two cases the cell had already shrunk to a point at which the membrane could no longer be distinguished from the middle region. Accordingly, for this case the confidence interval's z value was adjusted to 4.303.

Protein aggregates were most clearly visible after 6 hours, and were therefore only analysed at $t = 6$

In order to determine whether there was a statistically significant difference between calculated colocalisation, either at two points in time or when using different ROI selection methods (i.e. 2D vs 3D), we first used the f-test to determine the equality of the two associated variances and then used this result to select the appropriate two-sample, two-tailed t-test. The p-value obtained from this t-test could then be used to determine statistically significant differences ($p < 0.05$). In all cases the null hypothesis (H_o) was that the two calculated colocalisation metrics were the same.

4.3 Results

Figures 4.4-4.8 visualise representative cells for each considered point in time using both a 2D MIP and 45° 3D view. Note that a single 3D view cannot adequately convey a

true perception of the cell structure as perceived in virtual reality. In order to conserve space, only the view that generally best highlights the strength of the 3D approach is shown. For a better understanding of what is perceived in VR, the reader is referred to <https://youtu.be/nC2xh19KLGg>.

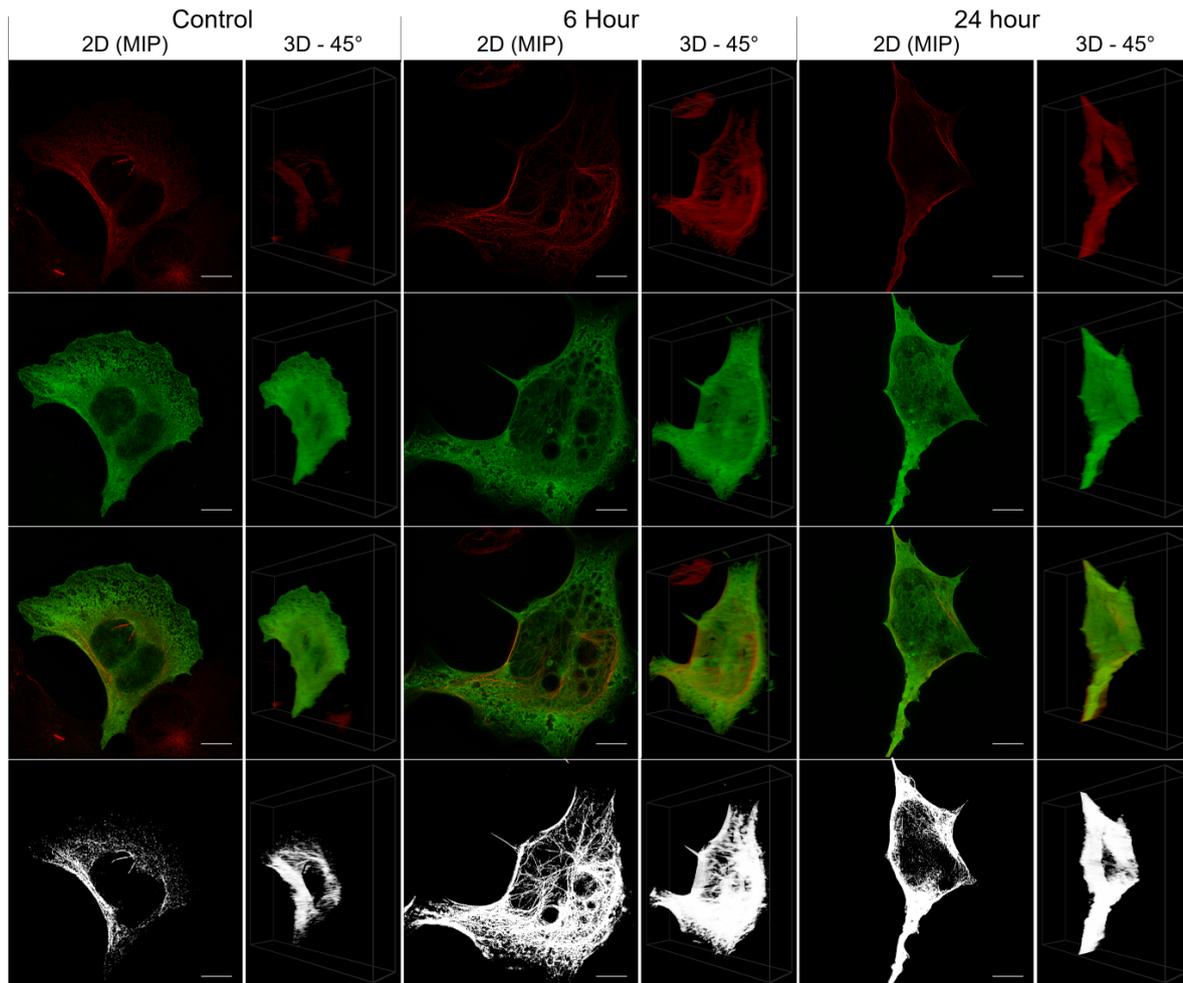


Figure 4.4: Showing the whole cell of representative cells for each considered time ($t = 0$, $t = 6$, $t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.

4.3.1 PCC and MOC results

Tables 4.2 and 4.3 show that, for all 3D ROI selections, there is a decrease in the colocalisation calculated using PCC and MOC at both $t = 6$ and $t = 24$ relative to the control, except for the microtubulin selection at $t = 6$ when using PCC. At $t = 6$, statistically significant differences relative to the control were observed for five of the nine selections (excluding the protein aggregates) for PCC (p-values < 0.032) and seven of the nine for MOC (p-values < 0.044). At $t = 24$, statistically significant differences relative to the control were observed for five of the nine selections for PCC (p-values < 0.038) and two of the nine for MOC (p-values < 0.049).

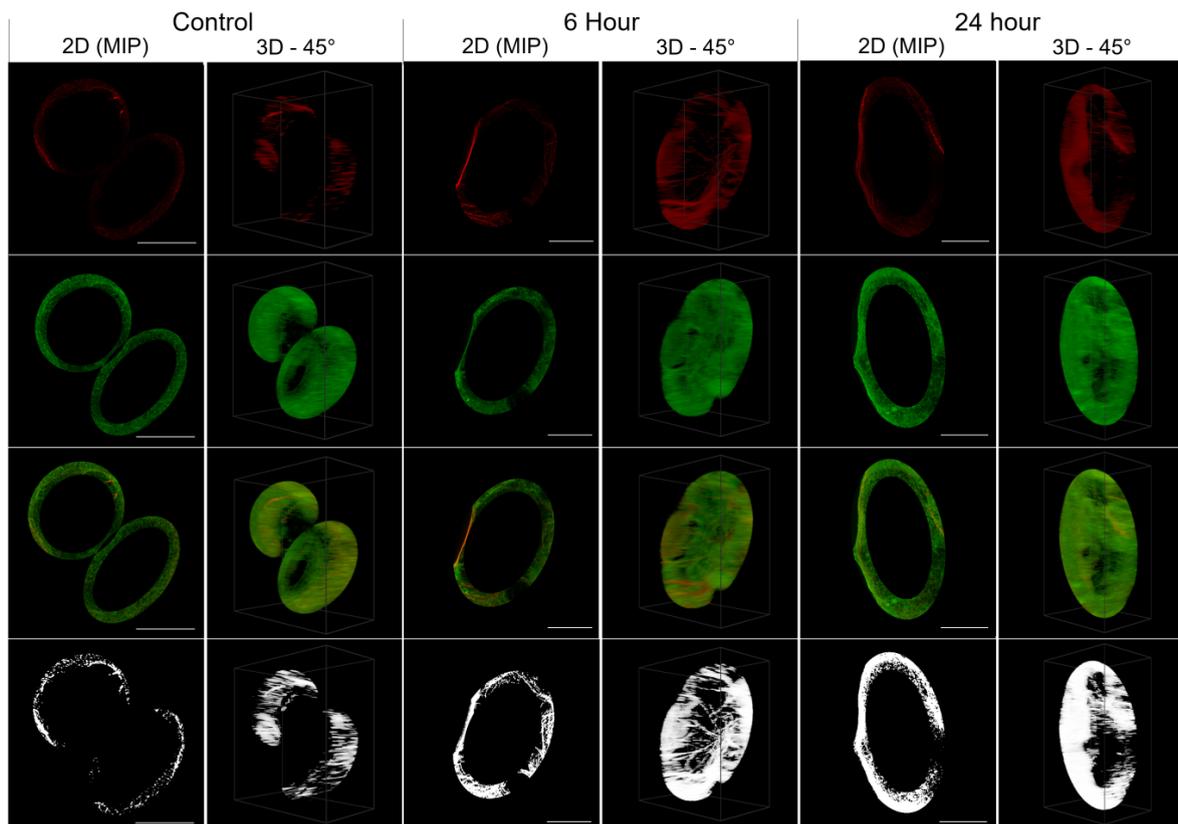


Figure 4.5: Showing the ellipsoidal perinuclear region of representative cells for each considered time ($t = 0$, $t = 6$, $t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.

Table 4.2 also shows that, when the colocalisation is calculated using a MIP ROI selection, there is an increase in PCC at both $t = 6$ and $t = 24$ relative to the control, except when selecting the whole cell at $t = 6$. On the other hand, Table 4.3 shows a decrease in MOC relative to the control for the same MIP selections for all regions from the control to $t = 6$ and control to $t = 24$, except for the microtubulin which showed an increase. None of these differences were statistically significant, however, except for the microtubulin PCC values at both $t = 6$ and $t = 24$ and the MOC for the middle and membrane regions at $t = 6$. No consistent change in PCC or MOC could be determined from $t = 6$ to $t = 24$, and none of these differences were statistically significant.

When considering the differences in PCC calculated using the 2D MIP and the 3D ROI selections, we find that there is no statistically significant difference between any of the selections at $t = 0$. For both $t = 6$ and $t = 24$, however, PCC indicated lower colocalisation for all 3D selections. At $t = 6$ and $t = 24$, respectively, eight of the ten (p-values < 0.022) and six of the nine (p-values < 0.003) differences were statistically significant. The MOC values were also lower when using the 3D ROIs than when using the MIP ROIs at all three considered points in time. At $t = 0$, the MOC of the 3D ROIs were statistically significantly lower than the MIP ROIs for three of the nine selections (p-values < 0.028). At $t = 6$ and $t = 24$, the MOC of the 3D ROIs were statistically significantly lower than the MIP ROIs for six of the ten (p-values < 0.046) and three of the nine (p-values < 0.05) selections, respectively.

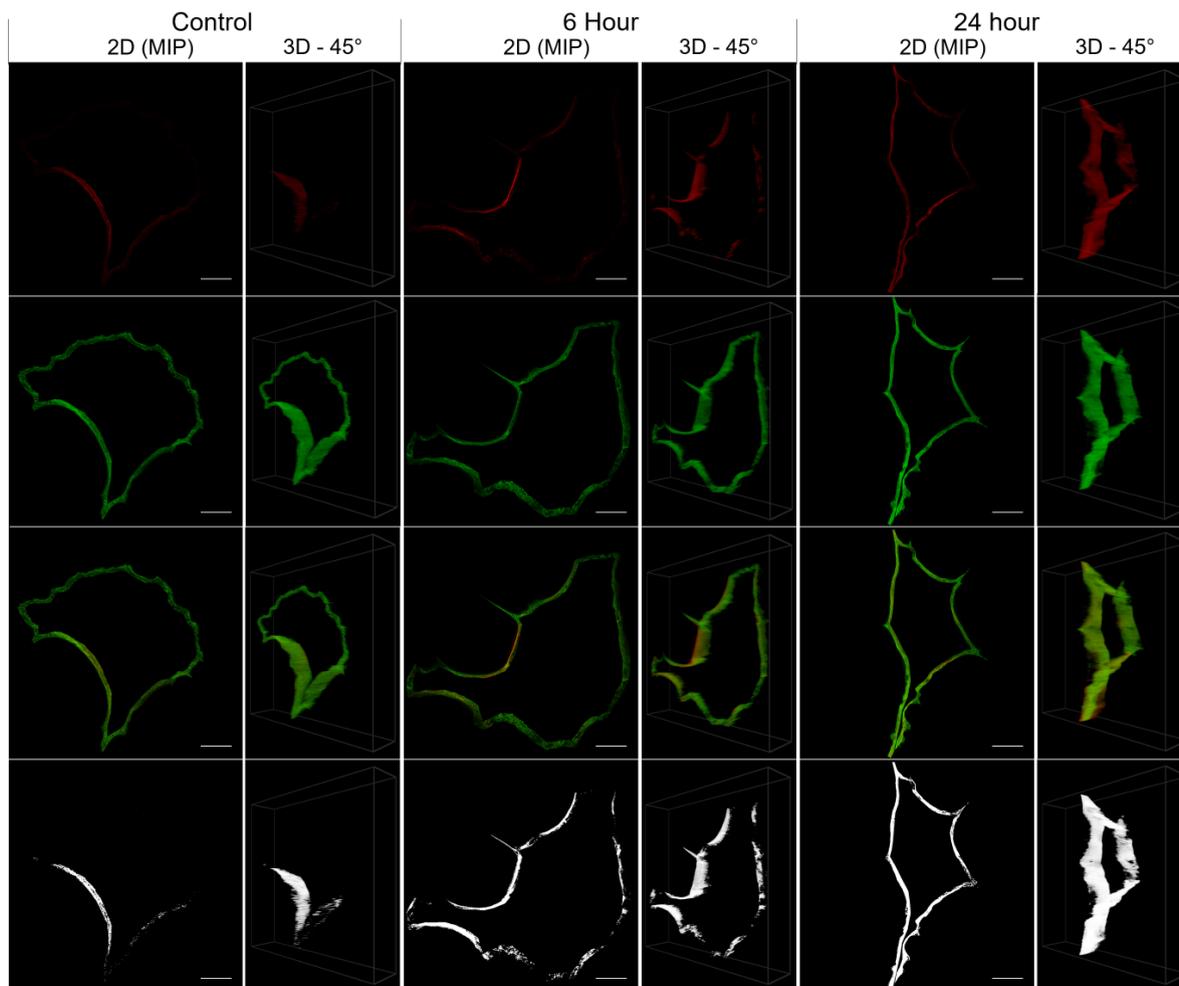


Figure 4.6: Showing the freehand membrane region of representative cells for each considered time ($t = 0$, $t = 6$, $t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.

4.3.2 MCC results

The results of the colocalisation analysis based on Manders' correlation coefficient (MCC) are compiled in Tables 4.4 and 4.5. We will refer to the MCC calculation for Tau simply as M1 and to the MCC calculation for acetylated α -tubulin simply as M2. When using 3D ROI selections, M1 showed an average increase in colocalisation with time for the whole cell, for the membrane and for the microtubulin. However, at $t = 24$ only the whole cell ($p = 0.023$) and the freehand membrane selection ($p = 0.026$) had increased significantly relative to the control. In contrast, M2 showed an average decrease in colocalisation from $t = 0$ through $t = 6$ to $t = 24$ for all 3D ROI selections except the cylindrical membrane selection at $t = 6$ and both membrane selections at $t = 24$. However, none of these decreases were statistically significant. When using MIP ROI selections, there was practically no variation in either M1 or M2, with most selections leading to a calculated MCC value of one. There were therefore no statistically significant differences between the values of M1 and M2 calculated between any of the three time instances $t = 0$, $t = 6$ and $t = 24$ when using the MIP-derived ROIs.

Of the M1 values shown in Table 4.4, 21 of the 28 MCC values were significantly

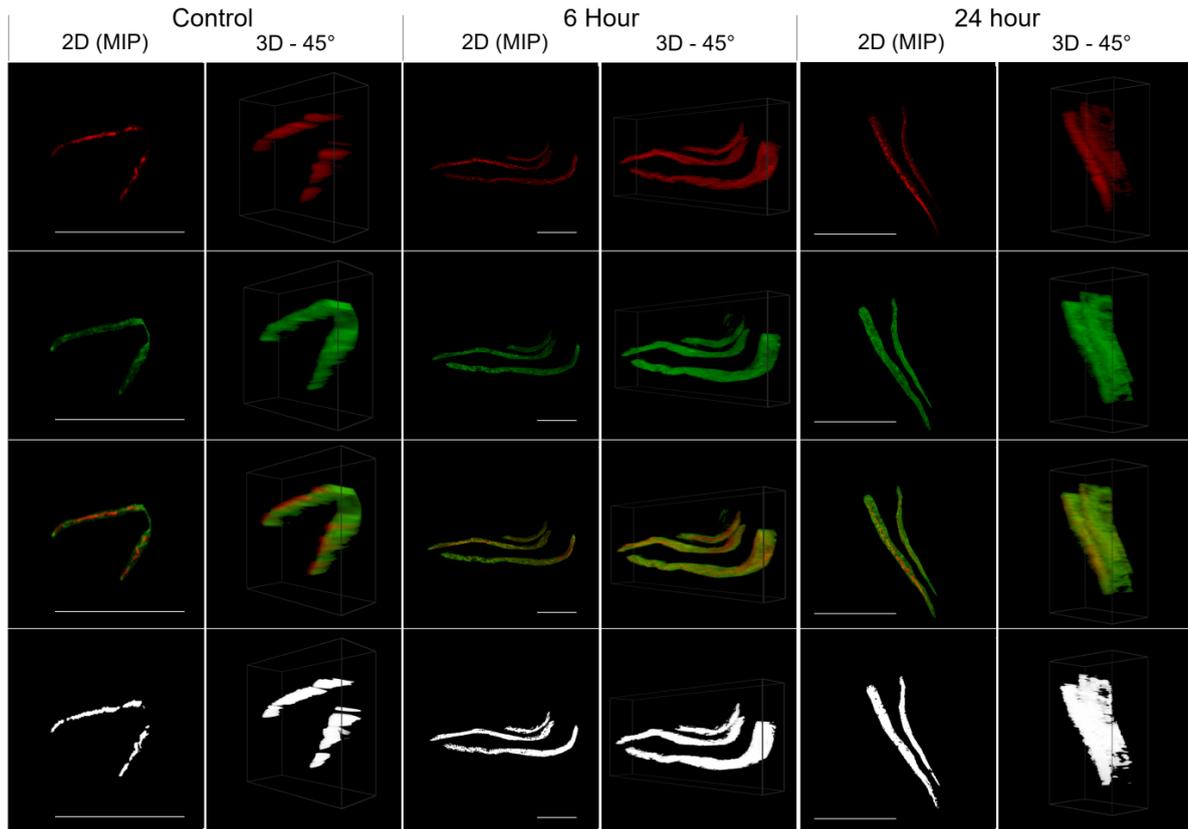


Figure 4.7: Showing some microtubulin strands of representative cells for each considered time ($t = 0$, $t = 6$, $t = 24$) with the maximum intensity projection on the left and the 45° 3D view on the right. Shown from top to bottom are the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels.

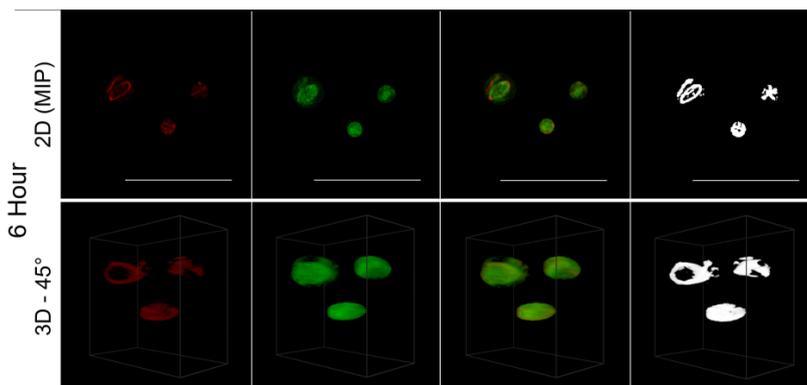


Figure 4.8: Showing some protein aggregates of representative cell for $t = 6$ with the maximum intensity projection at the top and the 45° 3D view with the ROI isolated from the whole cell at the bottom. Shown from left to right the acetylated α -tubulin (red), the Tau (green), the tubulin and Tau overlaid, and only the colocalised pixels/voxels. For the protein aggregates only the results at $t = 6$ were visualised and analysed.

higher (p-values < 0.046) when selecting ROIs in 2D than when selecting ROIs in 3D. For the M2 values shown in Table 4.5, 19 of the 28 MCC values were significantly higher (p-values < 0.047) in 2D compared to the 3D selections.

CHAPTER 4. IMPROVED ROI SELECTION AND COLOCALISATION ANALYSIS IN 3D FLUORESCENCE MICROSCOPY SAMPLES USING VIRTUAL REALITY 64

Table 4.2: Pearson's correlation coefficients calculated for Tau and acetylated α -tubulin in different areas of GT1-7 cells under control conditions and after exposure to chloroquine for 6 and 24 hours respectively. (Data are presented as mean \pm 95% confidence interval). * indicates significant differences relative to control ($p < 0.05$). A shaded block indicates that the ROI selection methods (i.e. 3D vs 2D) are statistically significantly different. In all cases, statistical significance was determined using the two-sample, two-tailed t-test. Protein aggregates were only analysed at the 6 hour time point.

	3D			2D (MIP)		
	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)
Whole Cell	0.72 \pm 0.07	0.60 \pm 0.07*	0.65 \pm 0.10	0.67 \pm 0.16	0.62 \pm 0.12	0.74 \pm 0.08
Perinuclear (Cylinder)	0.65 \pm 0.07	0.47 \pm 0.09*	0.47 \pm 0.12*	0.58 \pm 0.25	0.77 \pm 0.07	0.79 \pm 0.06
Perinuclear (Ellipsoid)	0.66 \pm 0.07	0.43 \pm 0.10*	0.45 \pm 0.14*			
Middle Region (Cylinder)	0.65 \pm 0.10	0.50 \pm 0.13	0.40 \pm 0.13*	0.55 \pm 0.31	0.79 \pm 0.05	0.80 \pm 0.10
Middle Region (Ellipsoid)	0.52 \pm 0.10	0.36 \pm 0.07*	0.31 \pm 0.19			
Membrane (Cylinder)	0.75 \pm 0.07	0.56 \pm 0.16*	0.48 \pm 0.24*	0.68 \pm 0.16	0.78 \pm 0.08	0.85 \pm 0.07
Membrane (Ellipsoid)	0.61 \pm 0.13	0.42 \pm 0.14	0.42 \pm 0.33			
Membrane (Freehand)	0.74 \pm 0.06	0.59 \pm 0.16	0.58 \pm 0.09*	0.69 \pm 0.12	0.73 \pm 0.10	0.80 \pm 0.05
Microtubulin	0.46 \pm 0.15	0.47 \pm 0.05	0.41 \pm 0.07	0.29 \pm 0.29	0.78 \pm 0.05*	0.82 \pm 0.05*
Protein Aggregates		0.25 \pm 0.16			0.82 \pm 0.04	

Table 4.3: Manders' overlap coefficient with procedure described in Table 4.2.

	3D			2D (MIP)		
	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)
Whole Cell	0.76 \pm 0.05	0.66 \pm 0.07*	0.67 \pm 0.10	0.79 \pm 0.08	0.76 \pm 0.06	0.77 \pm 0.07
Perinuclear (Cylinder)	0.80 \pm 0.04	0.65 \pm 0.10*	0.67 \pm 0.10*	0.86 \pm 0.05	0.78 \pm 0.08	0.80 \pm 0.06
Perinuclear (Ellipsoid)	0.81 \pm 0.04	0.65 \pm 0.07*	0.67 \pm 0.09*			
Middle Region (Cylinder)	0.81 \pm 0.05	0.69 \pm 0.08*	0.66 \pm 0.13	0.89 \pm 0.03	0.80 \pm 0.05*	0.80 \pm 0.10
Middle Region (Ellipsoid)	0.81 \pm 0.06	0.69 \pm 0.07*	0.66 \pm 0.13			
Membrane (Cylinder)	0.82 \pm 0.04	0.69 \pm 0.11*	0.71 \pm 0.22	0.88 \pm 0.03	0.79 \pm 0.08*	0.86 \pm 0.07
Membrane (Ellipsoid)	0.82 \pm 0.05	0.69 \pm 0.11*	0.73 \pm 0.21			
Membrane (Freehand)	0.80 \pm 0.02	0.68 \pm 0.14	0.72 \pm 0.09	0.84 \pm 0.05	0.75 \pm 0.09	0.81 \pm 0.05
Microtubulin	0.71 \pm 0.13	0.67 \pm 0.06	0.67 \pm 0.09	0.77 \pm 0.14	0.79 \pm 0.05	0.83 \pm 0.05
Protein Aggregates		0.69 \pm 0.03			0.82 \pm 0.04	

Table 4.4: Manders' correlation coefficient for Tau (M1) with procedure described in Table 4.2.

	3D			2D (MIP)		
	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)
Whole Cell	0.84 ± 0.09	0.91 ± 0.06	0.95 ± 0.03*	0.996 ± 0.004	0.998 ± 0.002	0.999 ± 0.001
Perinuclear (Cylinder)	0.96 ± 0.02	0.94 ± 0.02	0.96 ± 0.03	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00
Perinuclear (Ellipsoid)	0.97 ± 0.01	0.94 ± 0.03	0.96 ± 0.04			
Middle Region (Cylinder)	0.94 ± 0.04	0.93 ± 0.04	0.96 ± 0.02	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00
Middle Region (Ellipsoid)	0.97 ± 0.03	0.96 ± 0.03	0.97 ± 0.02			
Membrane (Cylinder)	0.85 ± 0.06	0.91 ± 0.04	0.96 ± 0.01	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00
Membrane (Ellipsoid)	0.91 ± 0.04	0.95 ± 0.04	0.97 ± 0.02			
Membrane (Freehand)	0.84 ± 0.09	0.91 ± 0.05	0.96 ± 0.03*	0.999 ± 0.001	1.00 ± 0.00	1.00 ± 0.00
Microtubulin	0.92 ± 0.10	0.93 ± 0.04	0.97 ± 0.03	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00
Protein Aggregates		0.98 ± 0.02			1.00 ± 0.00	

Table 4.5: Manders' correlation coefficient for acetylated α -tubulin (M2) with procedure described in Table 4.2.

	3D			2D (MIP)		
	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)	Control ($t = 0$)	6 Hour ($t = 6$)	24 Hour ($t = 24$)
Whole Cell	0.85 ± 0.04	0.82 ± 0.08	0.76 ± 0.13	0.92 ± 0.06	0.97 ± 0.02	0.77 ± 0.15
Perinuclear (Cylinder)	0.94 ± 0.04	0.88 ± 0.04	0.90 ± 0.04	0.999 ± 0.001	0.996 ± 0.004	0.998 ± 0.002
Perinuclear (Ellipsoid)	0.92 ± 0.06	0.85 ± 0.07	0.88 ± 0.04			
Middle Region (Cylinder)	0.94 ± 0.03	0.89 ± 0.07	0.88 ± 0.07	0.999 ± 0.001	0.998 ± 0.02	0.999 ± 0.001
Middle Region (Ellipsoid)	0.96 ± 0.04	0.91 ± 0.06	0.89 ± 0.07			
Membrane (Cylinder)	0.85 ± 0.06	0.87 ± 0.05	0.92 ± 0.10	0.99 ± 0.01	0.998 ± 0.002	1.00 ± 0.00
Membrane (Ellipsoid)	0.92 ± 0.06	0.90 ± 0.07	0.93 ± 0.09			
Membrane (Freehand)	0.86 ± 0.05	0.83 ± 0.05	0.90 ± 0.06	0.98 ± 0.01	0.99 ± 0.01	0.98 ± 0.02
Microtubulin	0.95 ± 0.08	0.91 ± 0.04	0.94 ± 0.03	1.00 ± 0.00	0.999 ± 0.001	0.999 ± 0.001
Protein Aggregates		0.92 ± 0.06			0.99 ± 0.01	

4.4 Discussion

When PCC and MOC were calculated using 3D ROI selections, all the values decreased after 6 and 24 hours relative to the control (except for the microtubulin at $t = 6$), with most of these decreases being statistically significant. These results support the notion that increased neuronal toxicity, i.e. autophagy dysfunction and protein aggregation, is

associated with a loss of colocalisation between Tau and tubulin, leading to a loss in microtubule stability. However, these decreases were not detected for PCC when using the 2D ROI selections. Instead, in this case, all selections showed average increases in colocalisation, with the microtubulin (Figure 4.7) showing a statistically significant increase (p-values < 0.026) at both $t = 6$ and $t = 24$ relative to the control. This is contrary to both the expected behaviour and the indication of the 3D analysis that there is no significant change in PCC for the microtubulin. The result of the 2D ROI selection is regarded as a false positive. Generally, 3D ROI selections must lead to a statistically more reliable estimate of colocalisation, since the entire sample is used in the calculation instead of just a projection. Since the MIP is based on maximum values along the projected dimension, it also typically leads to higher colocalisation estimates than when using the 3D ROI. This indicates a risk of false positives for 2D colocalisation analysis. In contrast to PCC, MOC mostly indicated an average decrease in colocalisation for both the 2D and 3D selections, where decreases observed for the middle and the cylindrical membrane regions were statistically significant. This seems to indicate that MOC is less affected by the limitations of a 2D projection than PCC.

For PCC, none of the differences between 2D and 3D ROI selections at the control ($t = 0$) were significant. This can be attributed to the initial (in the control condition) fairly good correlation between the colocalisation of Tau and MT, which leads to a projection that is a good approximation of the true 3D colocalisation. However, as soon as the cell is exposed to chloroquine and the neuronal toxicity increases, Tau and MT may start to dissociate in such a way that, when the volume is projected, PCC indicates a greater colocalisation than when considered in three-dimensional space. However, at $t = 0$ MOC and MCC were observed to be more sensitive than PCC and it was possible to determine some statistically significant differences between the 3D ROI selections when compared with the 2D MIP selections (Tables 4.3-4.5). Since the 3D analysis consistently showed these differences, it can be regarded as a more robust analysis approach than the 2D MIP analysis.

The strength of the 3D VR-enabled system is most clearly demonstrated in the case of the 3D perinuclear selection (Figure 4.5), which is one of the more complex structures in the cell that was considered. In this case we compared how ellipsoidal and cylindrical ROI selections influenced the differences between the colocalisation calculated at the control and at later time points. As was expected, the ellipsoidal selection, which is not possible in 2D, provided stronger evidence against the null hypothesis (H_0), which is that the two colocalisation metrics are the same, with smaller p-values (PCC $t = 6$: 0.005; PCC $t = 24$: 0.018; MOC $t = 6$: 0.004; MOC $t = 24$: 0.015) compared to the cylindrical selection (PCC $t = 6$: 0.009; PCC $t = 24$: 0.020; MOC $t = 6$: 0.014; MOC $t = 24$: 0.032). We believe that this is because more relevant data is taken into account when using the ellipsoidal ROI selection, especially at the top and bottom of the perinuclear region. A comparison of these selections is shown in Figure 4.2. In general, making selections in 3D increases the ability to detect differences in both PCC and MOC and we note specifically that no significant decrease was observed in PCC when using the traditionally considered MIP.

The analysis of MCC reveals that, when calculated using 2D MIP, many colocalisation estimates equal unity. When the MIP is generated, only the maximum intensities are projected and therefore the less-colocalised (lower) intensities are lost. This leads to the summation of the same intensities in the numerator and in the denominator of Equations 3 and 4 and a resultant MCC of unity. When making the same selections in 3D, all voxels, including those which are dimmer and considered less-colocalised, are taken into account,

leading to variations in the calculated MCC. This indicates that the 3D metric is more sensitive to variation in the data and that, by performing the analysis in 3D rather than in 2D, colocalisation in different regions can be compared.

For M1, an average increase was observed for most 2D and 3D ROI selections, with those for the whole cell (Figure 4.4) and the freehand membrane (Figure 4.6) selections at $t = 24$ being statistically significant (p-values < 0.026). This increase is in conflict with the decrease that was hypothesised and also calculated using the other metrics. We attribute this to an increase in the fluorescence brightness observed in the Tau fluorochrome at $t = 24$, which leads to a greater number of pixels having an intensity exceeding the predetermined colocalisation thresholds. This increase in fluorescence intensity of the fluorochromes after long-term exposure (i.e. 24 hours) to chloroquine is likely not due to an increase in mean Tau intensity, but rather due to toxicity-induced cell shrinkage. This shrinkage leads to an increased fluorochrome concentration within cell and thus an increase in brightness. This problem was specific to our setup and should not be taken as a general weakness of MCC.

It is worth pointing out that no significant difference between the colocalisation at $t = 6$ and $t = 24$ were found for any ROI selection for any of the considered metrics. From a biomedical perspective this is important, since only prolonged exposure to chloroquine led to significant cell death induction. This suggests that loss of colocalisation between Tau and MT may be an early event in neuronal injury, prior to the onset of cell death, which has not previously been shown. The dysfunction of Tau has, however, previously been considered to contribute to the collapse of the cytoskeleton [77]. This result, therefore, contributes to the understanding of the molecular pathology that unfolds in the pathogenesis, which sheds light on disease mechanisms, but may also be useful for unravelling novel therapeutic targets, although more research is needed.

When comparing all the results obtained for the 2D and 3D ROI selections we note that most are statistically significantly different (especially in the case of MCC). From this we conclude that the result of a colocalisation analysis is not only dependent on the metrics that are used, but also on the method of ROI definition. The results show that generally the 3D system is more sensitive than the 2D system and can decrease the occurrence of false positives and increase the sensitivity of the metrics.

Since 3D ROI selections include more data (voxels) than ROIs based on 2D MIPs, the resultant colocalisation metrics should deliver more robust results. Therefore, calculating a colocalisation estimate using an ROI defined in 3D offers greater confidence in the accuracy of the estimate than when it is defined in 2D on a MIP. For the same reason greater confidence is attributed to the estimates obtained using the freehand membrane selection over the cylindrical or ellipsoidal samples taken on the membrane, since the entire membrane is included in the calculation, instead of only samples along the membrane.

Using the 3D VR ROI selection system, specific regions could be analysed. These were the perinuclear region (Figure 4.5), the cell's middle regions (Figure 4.2), the cell membrane (Figures 4.2 and 4.6), the tubulin network (Figure 4.7) as well as protein aggregates (Figure 4.8). This 3D dissection of many cells allows greater precision of analysis. Since all selections and calculations remain in 3D, it leads to an increased resolving power of the colocalisation analysis. This stands in contrast to typically only analysing the whole-cell colocalisation using MIPs (Figure 4.4).

Using our 3D VR-assisted system, we were able to isolate very fine sub-cellular structures such as protein aggregates and microtubulin strands through precise ROI selection. This allowed colocalisation analysis to be considered specifically within those structures. This is not possible using conventional 2D analysis. These structures are visualised in

Figures 4.7 and 4.8.

Although it has previously been shown that chloroquine treatment enhances Tau aggregation [78], while wild type and phosphomutant Tau species do not aggregate in autophagy-deficient cells [79], the colocalisation profile under autophagy dysfunction has not previously been described in such precise and region specific manner. Here we describe a highly region-specific decrease in colocalisation upon autophagy disruption, which, to our knowledge, has not been described before. Similar results, albeit without region specific quantitative colocalisation analysis, have been reported in models of Tauopathy, with a distinct redistribution pattern of Tau and the microtubule network in relation to the point of dendritic degeneration [80]. This stresses the importance of being able to perform highly defined, region-specific colocalisation analysis, enabling the dissection of specific intracellular domains associated with neuronal pathology.

4.5 Conclusion

Despite the advances in fluorescence microscopy that allow detailed 3D data sets of cells to be generated, most colocalisation analysis is still performed using 2D maximum intensity projections. Our experiments indicate that enhanced and more sensitive colocalisation analysis is possible when carrying out three-dimensional region of interest selection. We focused on neuronal cells, specifically investigating the dissociation between Tau and the microtubulin network in an *in vitro* model of autophagy dysfunction, mimicking a key pathology in Alzheimer's disease, by analysing the change in colocalisation between Tau and acetylated α -tubulin at a control, 6 hour and 24 hour time points. After calculating several colocalisation metrics using both the 2D and 3D ROI selection methods, we compared the colocalisation estimated at the different points in time on the basis of the 2D and 3D ROIs. Statistically significant decreases in colocalisation over time were observed when using the 3D analysis, indicating that the dissociation of Tau and microtubulin may be an early event in neuronal injury prior to the onset of cell death. This tendency could not be consistently observed when using the 2D analysis. In fact, the 2D PCC analysis in some cases indicated increasing colocalisation with time. Generally, differences observed in the 2D analysis were statistically less significant when compared with the 3D analysis.

We were able to show how very fine structures, such as protein aggregates and microtubulin strands, can be isolated and analysed independently and with high precision. The work presented in this chapter is part of ongoing research in which the role of autophagy dysfunction in protein aggregation and tubulin stability in the context of neuronal toxicity is investigated, with the aim of discovering the key molecular mechanisms that govern neuronal fate. It is hoped that 3D virtual reality assisted ROI selection will allow the future analysis of samples of biological relevance can be interrogated with greater precision and control, thereby fully exploiting the potential of fluorescence-based image analysis in biomedical research.

In the next chapter we will move beyond the quantitative analysis of colocalisation, by means of metrics, to the qualitatively investigating of the underlying correlation between the fluorescence channels by means of a new spatial visualisation method.

4.6 Follow-up work

The method described in this chapter has since been applied to the investigation of the role of leucine-rich repeat kinase 2 (LRRK2) in mitochondrial function [81]. The results of this study have been published in the following article:

Neethling, A., Engelbrecht, L., Loos, B., Kinnear, C., Theart, R., Abrahams, S., Niesler, T., Mellick, G.D., Williams, M. and Bardien, S., 2019. Wild-type and mutant (G2019S) leucine-rich repeat kinase 2 (LRRK2) associate with subunits of the translocase of outer mitochondrial membrane (TOM) complex. *Experimental Cell Research*, 375(2), pp.72-79.

Chapter 5

Regression adjusted colocalisation colour mapping (RACC) for 3D qualitative colocalisation analysis

The content of this chapter was published as the following research article:

Theart, R.P., Loos, B. and Niesler, T.R., 2019. Regression adjusted colocalisation colour mapping (RACC): A novel biological visual analysis method for qualitative colocalisation analysis of 3D fluorescence micrographs. *PLOS one*, 14(11).

Up to now, we have mainly been concerned with the visualisation of fluorescence microscopy samples and the quantitative colocalisation analysis thereof. In this chapter, we expand our visualisation and analysis to include the qualitative investigation of colocalisation by means of a new visual analysis method that spatially visualises the correlation of the fluorescence channels.

5.1 Introduction

Fluorescence microscopy is a major driving force in modern biology and medicine, offering steadily increasing resolution and power of analysis. In such analyses, colocalisation, the geometric codistribution of two fluorescence channels (often referred to as signals), provides critical information indicating whether two proteins or structures of interest associate with one another. This is important for the understanding of biological processes and cellular functions. However, the objective is usually not merely to consider the spatial overlap of two fluorescence channels, since this would include coincidental overlap. Instead, it is of much greater importance to consider the correlation, or the proportional overlap, of two fluorescence channels within and between structures [74]. Therefore, for many colocalisation applications, it is desirable to accurately quantify the degree of colocalisation in the sample as well as to assess the location and intensity thereof clearly.

A common approach to quantifying colocalisation is the calculation of several colocalisation metrics, each of which highlights a particular aspect of the colocalisation and signal distribution throughout the sample or within an isolated region of interest (ROI). Some of the most notable and widely utilised among these metrics are the Pearson correlation coefficient (PCC), the Manders' Overlap coefficient (MOC) and the Manders' correlation

coefficient (MCC) [82] as presented in Section 3.1.1. These metrics calculate a single value that provides an indication of the overall correlation between the underlying colocalised fluorescence intensities over the analysis region as a whole. Although these measures are effective for the comparison of colocalisation between samples, especially when coupled with ROI selection, they are less suitable to convey any spatial information. Therefore, since sample investigations often require an understanding of how a fluorescence signal distributes throughout intracellular regions, another frequent approach to the analysis of colocalisation is by means of visualisation. Often this is achieved by overlaying the two fluorescence channel images and observing regions of overlap. For example, in the case of a red and green channel combination, the overlapping regions will be visualised in yellow. Although this approach provides a rapid overview of potentially colocalised signals, the ability to observe such yellow areas is highly dependent on the relative signal intensity of each channel. This is problematic since the intensity dynamics are rarely similar across different samples acquired through fluorescence microscopy.

Another common approach in the life sciences is to show the overlay of the fluorescence intensities together with a binary mask of the colocalised signal distribution. This binary mask is either shown by itself or superimposed on the fluorescence intensities as a single colour (often white) [83] as was described in Section 3.1.3. In this visualisation approach only the location of the colocalisation is shown. Limited or no indication is provided of the underlying intensities resulting in the observed colocalisation, or of the extent of the correlation between the channels. Lastly, visualisation of spatial colocalisation is most often performed two-dimensionally (2D) and only limited work has been undertaken to allow visualisation in three-dimensional (3D) space [57, 84].

In the rest of this chapter, we aim to address the above challenges, especially the limitations associated with showing the colocalised voxels only as a binary mask, by using a new approach that models the correlation of the underlying colocalised fluorescence signals spatially and visualises its distribution in three dimensions. With this newly proposed biological visual analysis method, which we refer to as *regression adjusted colocalisation colour mapping* (RACC), we aim to improve the 3D spatial interpretation of the colocalisation signal distribution within a sample in a robust way by capturing both the underlying channel intensities as well as their correlation. We demonstrate how this enhances the visualisation of colocalisation by analysing both synthetic data as well as biological samples, acquired through confocal as well as super-resolution techniques, with a focus on vesicle as well as tubulin network interactions. We further highlight the analytical strengths of RACC by integrating it with the virtual reality enabled 3D ROI selection tools developed in Chapter 4. RACC is intended as a complementary approach to existing colocalisation metrics. It has the specific aim of contributing an understanding of how the correlation between the underlying fluorescence intensities varies spatially.

5.1.1 Existing colocalisation intensity models

Several existing approaches consider the spatial quantification of the colocalisation in a sample, with each designed to highlight a certain aspect of the colocalisation. In order to contextualise the derivation of RACC, these models will be discussed briefly.

With the aim of improving the identification of colocalised structures at a sub-cellular level, enhanced colocalisation visualisation was pursued using a specifically designed dual-channel look-up-table (LUT) that maps fluorescence channel visualisations from Texas Red to magenta and FITC to cyan [72]. Red was used to indicate colocalised voxels that have a greater intensity in the Texas Red channel, green was used to indicate colocalised

voxels that have greater intensities in the FITC channel and yellow was used when both intensities were similar. In this way the relative intensities of the fluorescence channels in the colocalised voxel could be better discriminated [83].

Building on this work, two subsequent colocalisation visualisation methods were proposed, each having the advantage of not relying on the balanced staining of cells to ensure similar fluorescent signal intensities in both channels [85]. These methods were named the covariance method, which is a spatial representation of the PCC, and the multiply method, which is a spatial representation of MOC. Pixels having the greatest influence on the metrics are identified by visualising only those falling within the 99th percentile. In a similar way, several colocalisation metrics as well as the product of the PCC and MOC spatial maps can be visualised in what is referred to as a mixed map [86]. This mixed map is then used in an iterative classification process to generate a map that is a better approximation of the true colocalisation, as opposed to the coincidental colocalisation, thereby reducing the visualisation of false positives.

A currently very prominent approach to the visualisation of colocalisation provides a spatial representation of the correlation between two fluorescent signals [87], similar to the spatial representation of the PCC described above. For each pixel in the sample, a quantity termed the *normalised mean deviation product* (nMDP) is calculated as follows:

$$\text{nMDP}_{p_i} = \frac{(x_i - \bar{x})(y_i - \bar{y})}{(x_{p_{\max}} - \bar{x})(y_{p_{\max}} - \bar{y})} \quad (5.1)$$

with x_i and y_i representing the fluorescence intensities of the two fluorescence channels for the i^{th} pixel in the sample image, \bar{x} and \bar{y} representing the mean channel intensities and x_{\max} and y_{\max} representing the maximum channel intensities.

To calculate the mean and maximum channel intensities required in Equation 5.1, a Sobel filter is applied to each fluorescence channel of the image in order to make a region of interest (ROI) selection. In this way background intensities are removed independently for each fluorescence channel of the image. The use of a Sobel filter for background subtraction is however not required for nMDP and the authors note that other methods, such as manual intensity thresholding, are equally applicable. In the numerator of Equation 5.1, the product of the deviations from the mean within the ROI is calculated for fluorescence channels x and y . This product is then normalised by the product of the deviations of the maximum intensities from the respective means over the entire image. Using this equation, a new pseudo-colour image is generated that represents the degree to which colocalisation or non-colocalisation occurs at each pixel. The nMDP value can be either positive or negative, based on the intensity of the pixel relative to the mean. Positive values indicate that the pixel is colocalised and negative values indicate that it is not.

All the above methods aim to improve the spatial analysis of colocalisation by augmenting the standard visualisation of the overlap between two fluorescence channels in their original colours. Some of these attempt to introduce a visual representation of the correlation between the channels into the sample visualisation, by assigning a colour to each voxel based on its individual contribution to the PCC. Since the colocalisation metrics were not designed to be visualised directly, this can however introduce inconsistencies into the interpretation of the visualisation, as we will demonstrate. We also demonstrate how these inconsistencies are avoided by RACC. We will employ RACC to explore synthetic data as well as biological samples acquired through super resolution structured illumination microscopy (SR-SIM) and demonstrate the advantages of analysing colocalisation in

this manner. Since nMDP is most closely related to RACC in design and purpose, we will use this technique as a reference throughout.

5.2 Methods

When investigating colocalisation, usually the correlation and not merely the co-occurrence of fluorescence channel intensities is of interest. The intensity of a fluorescence channel in an image is dependent mainly on the abundance of the fluorochrome in that region of the cell. Therefore, when two proteins or structures of interest associate with one another, their fluorescent signals will usually have similar intensities in the same voxels of the image. One of the main aims of a colocalisation model is to spatially visualise the correlation between the two fluorescence channel intensities, specifically the degree to which colocalised voxels are *positively* correlated.

The two fluorescence intensities, for each voxel, can be represented by two random variables X and Y . In current micrographs, X and Y are usually discrete, with integer values between 0 and 255. For a given voxel i in the sample, this pair of intensities will henceforth be referred to as the *colocalisation intensity* \mathbf{q}_i and defined as:

$$\mathbf{q}_i = (x_i, y_i) \quad (5.2)$$

For a spatial visualisation of the correlation between fluorescence channels, the colocalisation intensity \mathbf{q}_i of each voxel should be assigned a colourmap value. This corresponds to a mapping from a 2D to a 1D subspace. Most existing methods, including nMDP (Equation 5.1), achieve this by a direct spatial visualisation of the PCC.

5.2.1 Direct visualisation of the PCC

The PCC can be estimated for X and Y , over N intensity bearing voxels in the ROI, using the sample correlation coefficient, as follows:

$$PCC = \hat{\rho}_{X,Y} = \frac{\sum_{i=1}^N (x_i - \bar{x}) \times (y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2 \times \sum_{i=1}^N (y_i - \bar{y})^2}} \quad (5.3)$$

In Equation 5.3, the denominator functions as a normalisation factor. The individual terms of the summation in the numerator can be viewed as an indication of how much the colocalisation intensity \mathbf{q}_i of each voxel in the ROI contributes to the estimated correlation coefficient $\hat{\rho}_{X,Y}$. Hence, some previous work visualised the normalised contribution of each voxel to $\hat{\rho}_{X,Y}$ directly using a colourmap [85, 86, 87]. This approach has some disadvantages, however.

Firstly, the numerator of Equation 5.3 can take either a positive or a negative value depending on the intensity of each fluorescence channel relative to the mean. To understand the implications of this, the relative contribution of all possible colocalisation intensities is illustrated in Figure 5.1A. The figure shows that the possible colocalisation intensities \mathbf{q}_i are divided into two positive and two negative quadrants about the means. The colour that nMDP assigns to each of these colocalisation intensities, using a colourmap, is shown in Figure 5.1B. Red indicates that the underlying voxel is labelled as colocalised while blue indicates that it is labelled as not-colocalised. This could lead to some discontinuities in the label that is assigned around the means, as is illustrated in Figure 5.1C for a colocalisation intensity distribution commonly associated with two

strongly correlated fluorescence channels [83]. While some follow-up studies using nMDP addressed this discontinuity at the means by using a more gradual transition between the colocalised and not-colocalised labels [88, 89, 90, 91, 92], these changes, however, alter the interpretation of the visualisation proposed by the authors of nMDP, since there is no longer a clear indication of which regions are colocalised and which are not. Furthermore, there is no straightforward way to redefine the colourmap, especially its start and end value, since nMDP can result in a wide range of numerical values.

A second disadvantage is that voxels with an intensity \mathbf{q}_i close to one of the means contribute least to the PCC value while voxels for which \mathbf{q}_i is far from the means contribute most. However, since only the distance from the mean and not the magnitude of the intensity is taken into account, voxels with colocalisation intensities at both the upper and lower extremes (indicated by * in Figure 5.1B) will be visualised with the same colour, resulting in an ambiguous visualisation. Furthermore, voxels for which \mathbf{q}_i is close to one of the means will be visualised with very low intensities by the colourmap. This suppresses the majority of the voxels, which are generally close to the means, thereby possibly losing critical information in the visualisation.

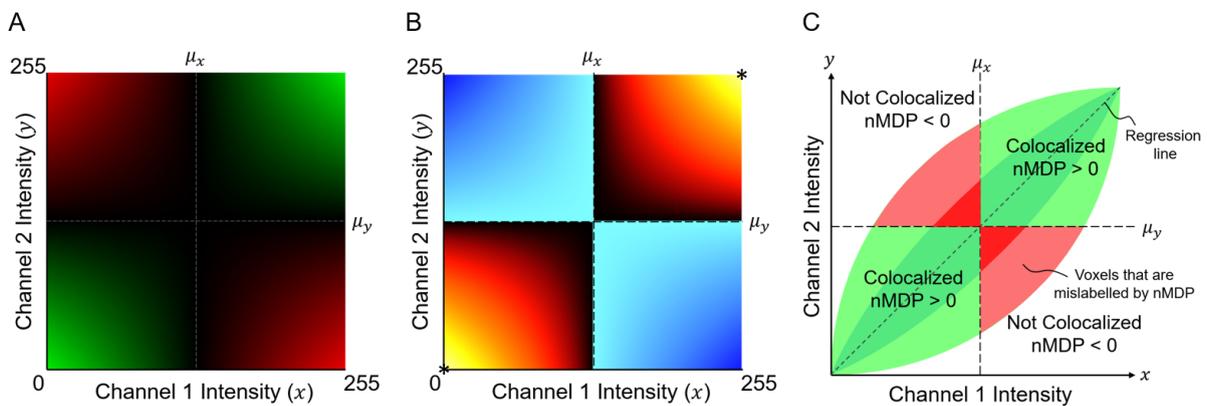


Figure 5.1: The direct visualisation of the PCC. The relative contribution of each voxel's colocalisation intensity \mathbf{q}_i to the correlation coefficient is determined by the individual terms of the summation in the numerator of Equation 5.3. With x and y axes representing the respective intensities of the two fluorescence channels, we assume a maximum intensity of 255, a mean intensity of 127 and a minimum intensity of 0 for each channel. A) The relative contribution to the PCC is shown for all possible values of \mathbf{q}_i , with green representing a positive and red representing a negative contribution. B) The colour that would be assigned to each voxel with intensity \mathbf{q}_i , using the nMDP colourmap [87], calculated with Equation 5.1. Note that, at the means, there are sudden changes in the colour assigned by nMDP. C) How nMDP labels each colocalisation intensity \mathbf{q}_i as colocalised or not-colocalised based on its position relative the means. A colocalisation intensity distribution commonly seen for strongly correlated fluorescence channels is shown, with darker shades indicating strong positive correlation and lighter shades indicating weaker positive correlation [83]. The areas marked in red correspond to positively correlated colocalisation that are indicated as negatively correlated when using nMDP.

Therefore, using the PCC formula directly to visualise the degree of correlation between the fluorescence channels has some critical drawbacks when analysing fluorescence micrographs. This is, however, not the design purpose of PCC and therefore it remains a useful way to quantify, with a single value, the degree to which the variation in intensity

between the two channels can be modelled by a simple linear relationship [74]. We will use this insight to develop RACC.

5.2.2 Visualising correlation using a regression line

Based on the limitations of existing visualisation methods discussed in the previous section, we develop RACC with three main objectives. Firstly, positively correlated colocalisation must be emphasised, while negatively correlated and coincidentally colocalised voxels should be suppressed. Secondly, regions of colocalised voxels with greater combined fluorescence intensities should also be emphasised, since greater fluorescence intensity is an indication of greater molecular density in the structure under investigation, and will assist in identifying molecular structures that strongly associate with one another. Thirdly, coincidentally colocalised voxels should be suppressed in the visualisation to reduce false positives.

In order to achieve these objectives we start by describing the linear relationship between the two fluorescence channel intensities. It is however not appropriate to use simple linear regression for this task, as has previously been suggested [83], since the accompanying assumptions of the ordinary least squares fit do not generally hold for colocalisation intensities [93]. Specifically, the assumption that one variable (the dependent variable) can be predicted from the other (the independent variable), and hence that the independent variable is a fixed, known constant, does not hold. Instead, both fluorescence intensities are sampled and are prone to observation errors. Therefore, we will use Deming regression, a special case of total least squares, that accounts for errors in the observations of both channel intensities [94].

For Deming regression, the regression line, or line of best fit, is described by:

$$y = \beta_0 + \beta_1 x \quad (5.4)$$

where x and y are the true values of the two fluorescence channel intensities and where β_0 and β_1 are the intercept and slope of the regression line, respectively. However, since the measured intensities \hat{x} and \hat{y} , have associated measurement errors ε_i and η_i , the true values are related to the measurements by:

$$\hat{x}_i = x_i + \eta_i, \quad \hat{y}_i = y_i + \varepsilon_i \quad (5.5)$$

In Deming regression it is assumed that the ratio of the variances of these two errors is known, and is defined as:

$$\lambda = \frac{\sigma_\eta^2}{\sigma_\varepsilon^2} \quad (5.6)$$

Since the measurements of the two channel intensities are usually made under the same conditions, we assume that the variances of the two errors are similar, and hence that $\sigma_\eta^2 \approx \sigma_\varepsilon^2$ and $\lambda \approx 1$. If, however, different acquisition parameters are used to generate the fluorescence images, the true ratio of the variances of the measurement errors must be used. For the remainder of the analysis we will assume $\lambda = 1$, which corresponds to a special case of Deming regression known as orthogonal regression [95].

In order to estimate β_1 we first calculate the sample mean and covariance of the random variables X and Y :

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i, \quad \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i \quad (5.7)$$

$$s_{xx} = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2, \quad s_{yy} = \frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})^2 \quad (5.8)$$

$$s_{xy} = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y}) \quad (5.9)$$

Now β_0 and β_1 are estimated as follows [94]:

$$\hat{\beta}_1 = \frac{\lambda s_{yy} - s_{xx} \pm \sqrt{(\lambda s_{yy} - s_{xx})^2 + 4\lambda s_{xy}^2}}{2\lambda s_{xy}} \quad (5.10)$$

$$\hat{\beta}_0 = \bar{y} - \hat{\beta}_1 \bar{x} \quad (5.11)$$

Because in general $\hat{\beta}_1$ has two solutions, the solution with the same sign as the covariance s_{xy} is conventionally chosen [96]. If the fluorescence channel intensities are positively correlated, the covariance, and therefore $\hat{\beta}_1$, will also be positive. However, in some rare cases the covariance can be negative, such as when the fluorescence channel intensities are not correlated or the colocalised intensities are very sparse. Our aim is to highlight voxels for which both fluorescence intensities are high, and hence we must ensure a positive slope for the regression line. Therefore, in our case, the positive solution for $\hat{\beta}_1$ will always be chosen.

It is standard accepted procedure in fluorescence based colocalisation analysis to remove background noise by applying thresholds to both fluorescence channels. We will refer to these thresholds as T_{ch1} and T_{ch2} . In the experiments we present later, we manually determine these thresholds. However, any method of determining these thresholds may be used, for example automated thresholding [84]. A voxel is only considered colocalised if both channel intensities are above their respective thresholds. Therefore, the means and covariances in Equations 5.7-5.9 are only calculated using the colocalisation intensities \mathbf{q}_i above these thresholds.

We will consider the regression line (Equation 5.4) as a 1D subspace onto which the colocalisation intensities \mathbf{q}_i are mapped, as illustrated in Figure 5.2. In order to perform this mapping, it is helpful to represent the regression line as a vector \mathbf{p} , that passes through two points \mathbf{p}_0 and \mathbf{p}_1 , representing the two extremes of the 1D subspace.

The regression line must intersect one of the two lines defined by the channel thresholds. Figure 5.2 shows that this intersection \mathbf{p}_0 can be calculated as follows:

$$\mathbf{p}_0 = (x_{p_0}, y_{p_0}) = \begin{cases} (T_{ch1}, T_{ch1}\hat{\beta}_1 + \hat{\beta}_0), & \text{if } T_{ch2} \leq T_{ch1}\hat{\beta}_1 + \hat{\beta}_0 \\ (\frac{T_{ch2} - \hat{\beta}_0}{\hat{\beta}_1}, T_{ch2}), & \text{if } T_{ch2} > T_{ch1}\hat{\beta}_1 + \hat{\beta}_0 \end{cases} \quad (5.12)$$

Similarly, \mathbf{p}_1 denotes the intersection between the regression line and the maximum possible fluorescence intensity I_{\max} , and is calculated as follows:

$$\mathbf{p}_1 = (x_{p_1}, y_{p_1}) = \begin{cases} (\frac{I_{\max} - \hat{\beta}_0}{\hat{\beta}_1}, I_{\max}), & \text{if } \hat{\beta}_0 \geq I_{\max}(1 - \hat{\beta}_1) \\ (I_{\max}, I_{\max}\hat{\beta}_1 + \hat{\beta}_0), & \text{if } \hat{\beta}_0 < I_{\max}(1 - \hat{\beta}_1) \end{cases} \quad (5.13)$$

The colocalisation intensity of each voxel \mathbf{q}_i can now be projected perpendicularly onto the regression line \mathbf{p} , resulting in a point \mathbf{p}_i . To achieve this projection we define

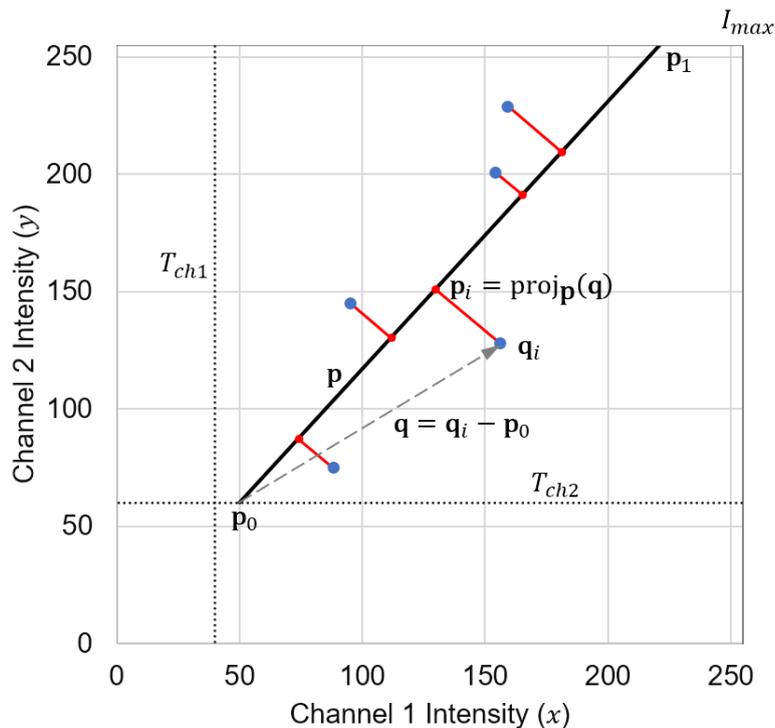


Figure 5.2: The projection of colocalised voxels onto the colourmap. A colocalisation intensity \mathbf{q}_i is projected to a 1D subspace defined by \mathbf{p} . A projection onto point \mathbf{p}_0 corresponds to a colourmap value of 0, a projection onto point \mathbf{p}_1 corresponds to a colourmap value of 1, and intermediate projections correspond to a colourmap value varying linearly between these values.

the vector from \mathbf{p}_0 to \mathbf{q}_i as \mathbf{q} (Figure 5.2) from which it follows that:

$$\mathbf{p}_i = (x_{p_i}, y_{p_i}) = \text{proj}_{\mathbf{p}}(\mathbf{q}) = \frac{\mathbf{q} \cdot \mathbf{p}}{\mathbf{p} \cdot \mathbf{p}} \mathbf{p} + \mathbf{p}_0 \quad (5.14)$$

This projection is used to assign a colourmap value C_i between 0 and 1 based on the linear position of \mathbf{p}_i along \mathbf{p} . This value is in turn used to assign a colour to the i^{th} voxel by means of a colourmap. Note that if \mathbf{p}_1 does not correspond exactly to the intensity (I_{\max}, I_{\max}) , the perpendicular projection of the colocalisation intensity \mathbf{q}_i may result in some points \mathbf{p}_i positioned beyond the maximum of the colourmap \mathbf{p}_1 . The colourmap value C_i of such voxels are assigned a colourmap value of 1. Similarly, if \mathbf{p}_0 does not correspond exactly to the intensity (T_{ch1}, T_{ch2}) , \mathbf{p}_i may lie below \mathbf{p}_0 . These voxels are assigned a colourmap value of 0. Therefore, C_i can be calculated as follows:

$$C_i = \begin{cases} 0, & \text{if } x_{p_i} \leq x_{p_0} \\ \frac{x_{p_i} - x_{p_0}}{x_{p_1} - x_{p_0}}, & \text{if } x_{p_0} < x_{p_i} < x_{p_1} \\ 1, & \text{if } x_{p_i} \geq x_{p_1} \end{cases} \quad (5.15)$$

To understand the effect of Equation 5.15 on the voxels in a sample, we plot the colourmap intensity C_i that would be used for every possible colocalisation intensity \mathbf{q}_i , for a particular regression line, in Figure 5.3A using the perceptually uniform Magma colourmap. From this it becomes clear that voxels that have high fluorescence intensities for both x and y are assigned a high value of C_i , while lower fluorescence intensities

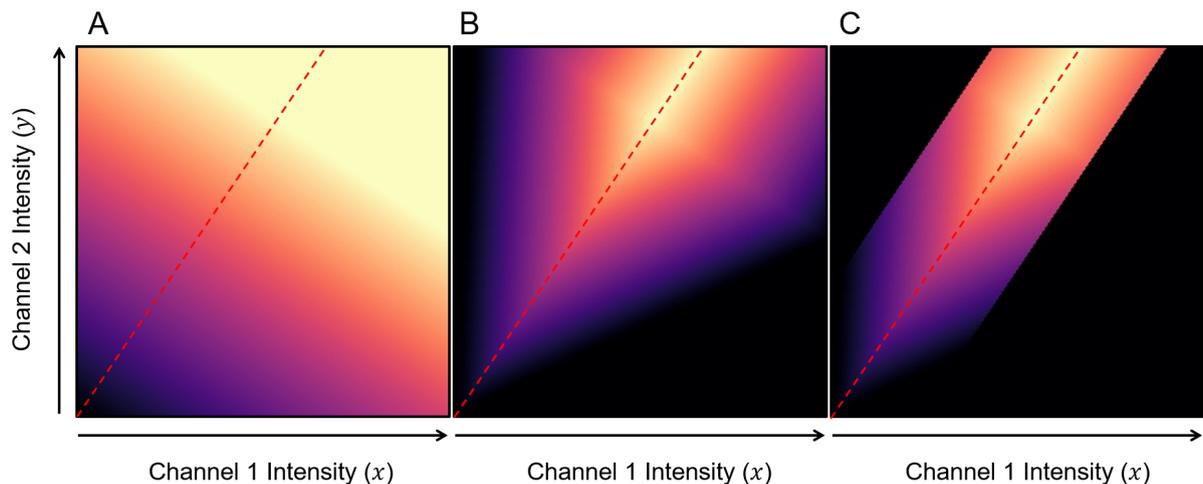


Figure 5.3: The effect of the colour mapping model on possible colocalisation intensities. Using Figure 5.2 as a reference, we consider a specific regression line, shown as a red dashed line, with $\hat{\beta}_1 = 1.5$ and $\hat{\beta}_0 = 0$ for illustration. We visualise the corresponding colourmap value (C_i) calculated for each possible colocalisation intensity $\mathbf{q}_i = (x_i, y_i)$ using the Magma colourmap, where black corresponds to $C_i = 0$ and bright yellow corresponds to $C_i = 1$. For clarity, both thresholds T_{ch1} and T_{ch2} are zero in this example. A: The colourmap values C_i using Equation 5.15. B: The colourmap values C'_i using Equation 5.17 with $\theta = 60^\circ$. C: The final colourmap values C_i using Equation 5.19 with $d_t = 0.2$ and $x_{p_{\max}} = 0.8I_{\max}$.

are assigned lower values of C_i . Equation 5.15 therefore satisfies our second objective of highlighting voxels with greater combined fluorescence intensities. However, it does not yet satisfy our first and third objectives, since the model does not suppress voxels that are not positively correlated. Instead, it assigns the same colourmap value to all colocalisation intensities that lie on a line perpendicular to the regression line.

To fulfil the first objective, we recall that the regression line represents correlation in the data and therefore we seek to suppress voxels for which \mathbf{q}_i is far from the regression line. To achieve this we calculate the normalised perpendicular distance d from \mathbf{q}_i to the regression line \mathbf{p} and adjust the colourmap value accordingly. This normalised distance d is given as:

$$d = \frac{|\mathbf{q}_i - \mathbf{p}_i|}{I_{\max}} = \frac{\sqrt{(x_i - x_{p_i})^2 + (y_i - y_{p_i})^2}}{I_{\max}} \quad (5.16)$$

The normalisation by I_{\max} ensures that $0 \leq d \leq 1$.

The colourmap value C_i originally calculated by Equation 5.15 is then attenuated linearly with the distance d . This is achieved by projecting each colocalisation intensity q_i onto the regression line at an angle θ to the perpendicular, as illustrated in Figure 5.4. The value of this *penalisation factor* θ therefore determines the extent to which the colourmap value C_i is attenuated as d increases. Attenuation is achieved using the following equation, with C_i calculated in Equation 5.15.

$$C'_i = \begin{cases} C_i - d \tan(\theta), & \text{if } C_i > d \tan(\theta) \\ 0, & \text{if } C_i \leq d \tan(\theta) \end{cases} \quad (5.17)$$

In Equation 5.17 it has been assumed that $0^\circ \leq \theta < 90^\circ$. A penalisation factor of 0° corresponds to no attenuation, and hence $C'_i \equiv C_i$, while factors close to but excluding 90° correspond to maximum attenuation, effectively making $C_i = 0$, and completely suppressing all voxels for which \mathbf{q}_i is not positioned precisely on the regression line. The effect of Equation 5.17 on all possible colocalisation intensities \mathbf{q}_i is visualised in Figure 5.3B for $\theta = 60^\circ$. The penalisation factor θ is the only adjustable parameter in the RACC method and can be determined interactively by varying it over the range $[0, 90)$. This allows the investigator to adjust the degree to which the voxels that show weak correlation between the fluorescence colours should be suppressed. Choosing a low penalisation factor (low value of θ) can cause the visualisation to become saturated with high colourmap values, while choosing a high penalisation factor (high value of θ) can lead to saturation with low colourmap values. The best choice will differ between samples. Once this value is found for a particular sample or condition, similar samples can subsequently be visualised with the same value to allow for the best and standardised comparison. Attempts can be made to automatically determine the value of the penalisation factor, however for the samples we analysed in this chapter we found that a value of $\theta = 45^\circ$ often produces favourable results.

To meet our third stated objective of suppressing coincidental colocalised voxels from the visualisation, we introduce a distance threshold d_t (Figure 5.4) which we use to remove outliers around the regression line. Only voxels for which the distance d between the colocalisation intensity \mathbf{q}_i and the regression line is below this threshold are visualised. The threshold d_t is determined automatically to include 99% of all voxels with intensities above the channel thresholds T_{ch1} and T_{ch2} . Assuming a normal distribution of the intensities, this corresponds to approximately three standard deviations from the mean. The effect of applying a distance threshold $d_t = 0.2$ is illustrated in Figure 5.3C. The corresponding colourmap value is calculated using Equation 5.18:

$$C''_i = \begin{cases} C'_i, & \text{if } d \leq d_t \\ 0, & \text{if } d > d_t \end{cases} \quad (5.18)$$

Finally, in scenarios where most voxels have low intensities for both channels, for example due to low intensity in the fluorescence microscopy acquisition parameters, there will be an ineffective utilisation of the colourmap spectrum and an associated limited variation in the visualisation. To address this, we define a point $\mathbf{p}_{\max} = (x_{p_{\max}}, y_{p_{\max}})$ on the regression line that indicates the maximum represented intensities for the image. Any projected point \mathbf{p}_i beyond \mathbf{p}_{\max} is assigned a colourmap value of 1. The point \mathbf{p}_{\max} is calculated to include 99% of all sample data above the channel thresholds. Since all colocalisation intensities are projected perpendicularly onto the regression line, the imposition of \mathbf{p}_{\max} has the effect of saturating the colourmap for voxels whose projected intensities fall beyond \mathbf{p}_{\max} , indicated in green in Figure 5.4. This makes more of the colourmap available for the representation of lower colocalisation intensities, thereby enabling clearer data interpretation. This re-scaling of the applied colourmap is achieved by replacing \mathbf{p}_1 with \mathbf{p}_{\max} in Equation 5.15. The researcher may, however, choose not to apply this scaling in which case \mathbf{p}_1 is retained.

Given \mathbf{q}_i and θ and obtaining the parameters \mathbf{p}_0 , \mathbf{p}_i and d from Equations 5.12, 5.14

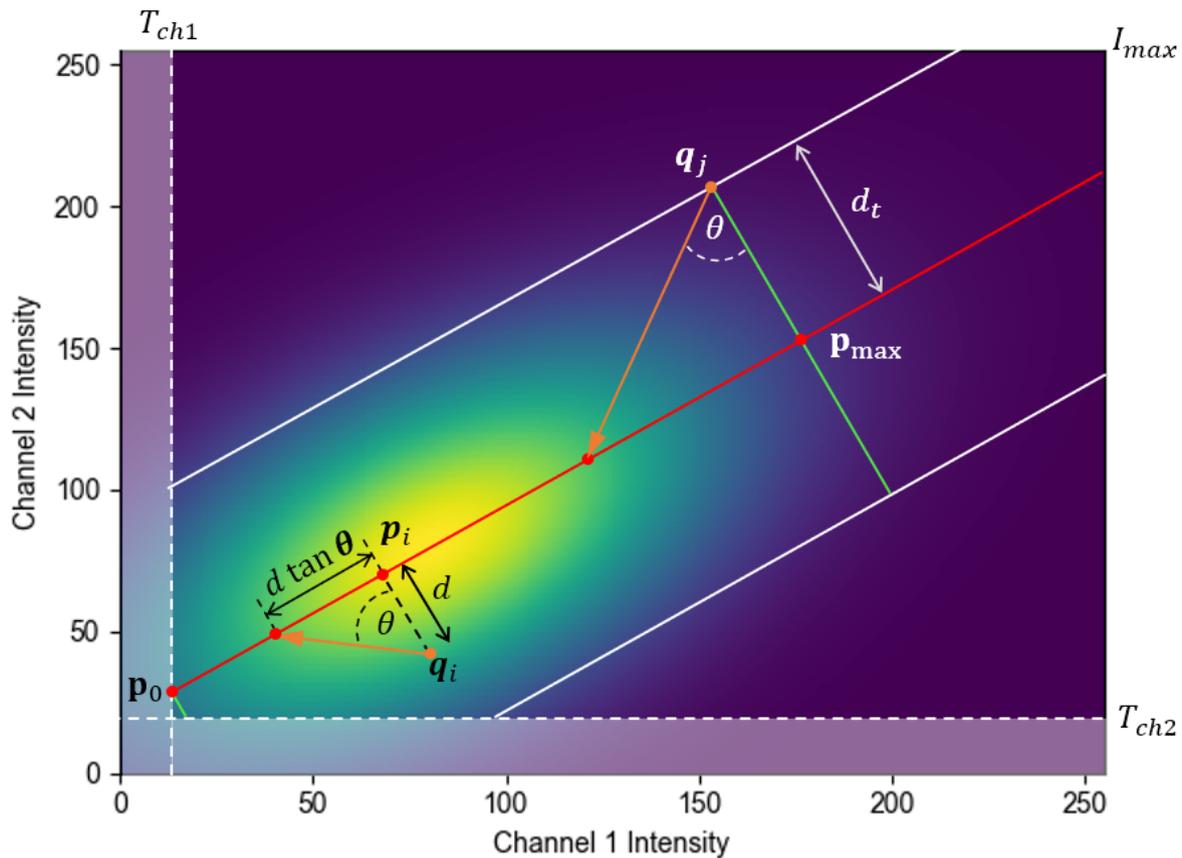


Figure 5.4: The illustration of the effect of the colourmap thresholds using a synthetically generated scatter plot, where the colourmap brightness indicates the frequency of voxels with a given colocalisation intensity. The regression line is shown in red, with two example intensity pairs, \mathbf{q}_i and \mathbf{q}_j , shown in orange. A penalisation factor θ is applied when projecting both \mathbf{q}_i and \mathbf{q}_j onto the regression line. Only the sample data above the channel thresholds, T_{ch1} and T_{ch2} , are considered when determining the distance threshold d_t , shown in white. Any colocalisation intensity \mathbf{q}_i beyond the maximum point \mathbf{p}_{max} is assigned a colourmap value of 1. Similarly any \mathbf{q}_i below the minimum point \mathbf{p}_0 is assigned a colourmap value of 0. These thresholds are indicated with two green lines at the respective points. Both d_t and \mathbf{p}_{max} are determined independently to include 99% of the data, thereby excluding outliers from the visualisation. The final colourmap value C_i varies linearly between 0 to 1 from \mathbf{p}_0 to \mathbf{p}_{max} .

and 5.16, Equation 5.15 can now be rewritten in its final form as follows:

$$C_i = \begin{cases} 0, & \text{if } x_{p_i} \leq d(x_{p_{max}} - x_{p_0})\tan(\theta) + x_{p_0} \text{ or } d > d_t \\ \frac{x_{p_i} - x_{p_0}}{x_{p_{max}} - x_{p_0}} - d \tan(\theta), & \text{if } d(x_{p_{max}} - x_{p_0})\tan(\theta) + x_{p_0} < x_{p_i} < x_{p_{max}} \\ 1 - d \tan(\theta), & \text{if } x_{p_i} \geq x_{p_{max}} \end{cases} \quad (5.19)$$

The result of Equation 5.19 is a colourmap value C_i between 0 and 1 for any pair of measured fluorescence intensities $\mathbf{q}_i = (x_i, y_i)$.

5.3 Results and discussion

With the aim of improving the 3D spatial visualisation of colocalisation in biological samples by incorporating both the correlation and the underlying fluorescence channel intensity in the visualisation, we developed the regression adjusted colocalisation colour mapping (RACC).

In the following, we will apply and validate the RACC algorithm by visualising colocalisation in both synthetically generated data as well as in three distinct biological samples. This is intended to demonstrate the advantages offered by RACC.

In the first biological sample we visualise the colocalisation between α/β tubulin and acetylated tubulin. In the second and third samples we visualise the colocalisation between two organelles, lysosomes and autophagosomes, and between autophagosomes and the tubulin network, respectively, in order to investigate their fused state which is of biomedical importance. Due to the small size of the organelles and the fineness of the filamentous tubulin structures, it is often challenging to accurately investigate the degree of colocalisation between them. These structures were, therefore, deliberately chosen to demonstrate the performance of RACC in challenging scenarios. The synthetic data was designed to mimic and better demonstrate aspects of these biological samples.

We will compare RACC (Equation 5.19) to the nMDP (Equation 5.1) throughout, since nMDP is currently the most widely used method of spatially visualising the correlation of the colocalised voxels using a colourmap. For both RACC and nMDP, all parameters were calculated over the entire 3D z-stack. For purposes of comparison, we used manual intensity thresholding to remove the background intensities for both methods. These thresholds are shown in the scatter plots accompanying the figures. Part of the design of RACC is the automatic adjustment of the maximum colourmap value to ensure the optimal use of the available spectrum. For nMDP, however, it is standard practice to fix the minimum and maximum values to -1.0 and 1.0, respectively [87]. In our experience, this does not always generalise optimally, especially since nMDP can produce values beyond these limits. Therefore, to ensure the best visualisation for each image set, we adjusted the minimum and maximum values of the nMDP colourmap manually. These values are indicated on the nMDP colourbars in the figures. Similarly, the penalisation factor θ that was used for each sample is shown on the colourbar.

5.3.1 Visualising synthetic data

In order to better understand how RACC visualises colocalisation in biological samples we generated synthetic data that is illustrative of some aspects of the biological samples that we investigate. We consider three sets of synthetic data, shown in section I, II and III of Figure 5.5. Section I shows two perfectly overlapping cylinders of the same size. These model the overlap between α/β tubulin and acetylated tubulin that will be seen in the biological samples. Section II shows two partially overlapping spheres, and are an idealisation of autophagosome-lysosome fusion. Finally, section III shows a slightly overlapping cylinder and sphere, representing autophagosome-tubulin interaction. This scenario is similar to the start of autophagosome-lysosome fusion which we modelled using two slightly overlapping spherical structures. For all synthetic images, the highest colour intensity occurs at the centre of the spheres and cylinders and decreases towards the surface of these structures. This mirrors the observed fluorescence intensities in the biological samples. Furthermore, in the case of the synthetic data, it is often insightful to consider the maximum intensity projection (MIP), which provides an internal perspective

of the 3D data sets. This is useful because most of the variation in fluorescence intensity occurs inside the volume and not at the surface.

5.3.1.1 Technical discussion

Figure 5.5A shows the colocalisation between red and green channels as yellow, and Figure 5.5B shows the same colocalised voxels in white. These are the most common ways of visualising colocalisation. From Figure 5.5B it is expected that a colocalisation colourmap method should only assign a colourmap value to the voxels in white. This value must vary with the channel intensities of the colocalised voxels in a manner similar to that seen in Figure 5.5A.

When considering Figure 5.5C, we recall that nMDP visualises colocalised voxels in shades of red to yellow, while voxels that are not-colocalised are visualised in shades of blue (see Figure 5.1B). By considering sections II and III of Figure 5.5C, it is clear that the nMDP assigns a colourmap value to all voxels for which either fluorescence channel is present, and not only to those that would normally be considered colocalised (indicated as white in Figure 5.5B). Since it is common for the channel means (indicated by a red dot in Figure 5.5F) to be greater than the intersection point of the two channel thresholds, lower intensity voxels are not removed from the visualisation. This leads some of these voxels to be assigned colocalised colourmap values according to Figure 5.1B, which could misrepresent the data. Figure 5.5D on the other hand, shows that RACC only assigns colours to voxels that are considered colocalised and does not visualise voxels that are not colocalised. This occurs because RACC is only calculated for voxels for which *both* fluorescence intensities are above the channel thresholds, whereas nMDP is calculated when a single fluorescence channel is above the respective channel threshold. Figure 5.5E visualises this difference by showing regions that are labelled as colocalised by nMDP but not by RACC in yellow, and regions for which the converse is true in magenta. The same information is also presented quantitatively in Table 5.1 which shows the percentage of the total voxels that were classified differently by RACC and nMDP.

Table 5.1: Percentage of the total voxels classified differently by RACC and nMDP.

	Perfectly overlapping cylinders	Partially overlapping spheres	Slightly overlapping sphere-cylinder	α/β tubulin and acetylated tubulin	Autophagasomes and lysosomes	Autophagasomes and tubulin
RACC and nMDP match for colocalised	100.0%	8.0%	0.4%	0.4%	2.4%	2.3%
RACC colocalised, nMDP not colocalised	0.0%	10.4%	0.0%	0.5%	2.6%	3.4%
RACC not colocalised, nMDP colocalised	0.0%	49.6%	55.6%	66.6%	67.8%	61.9%
RACC and nMDP match for not colocalised	0.0%	32.1%	44.0%	32.6%	27.2%	32.4%

In the case of the perfectly overlapping cylinders (Figure 5.5 section I), both fluorescence channels are perfectly correlated. Initially the nMDP decreases from a high intensity at the centre of the cylinder to its periphery. However, at lower fluorescence intensities an increase is again observed, leading to an unintuitive interpretation of the correlation in the colocalisation. This increase is due to the nMDP calculation only considering the deviation from the mean and not the magnitude of the intensity, as shown in Figure 5.1B. RACC shows a linear decrease from the high fluorescence intensities at the centre of the cylinder to the lower intensities at its periphery, thereby highlighting volumes with higher

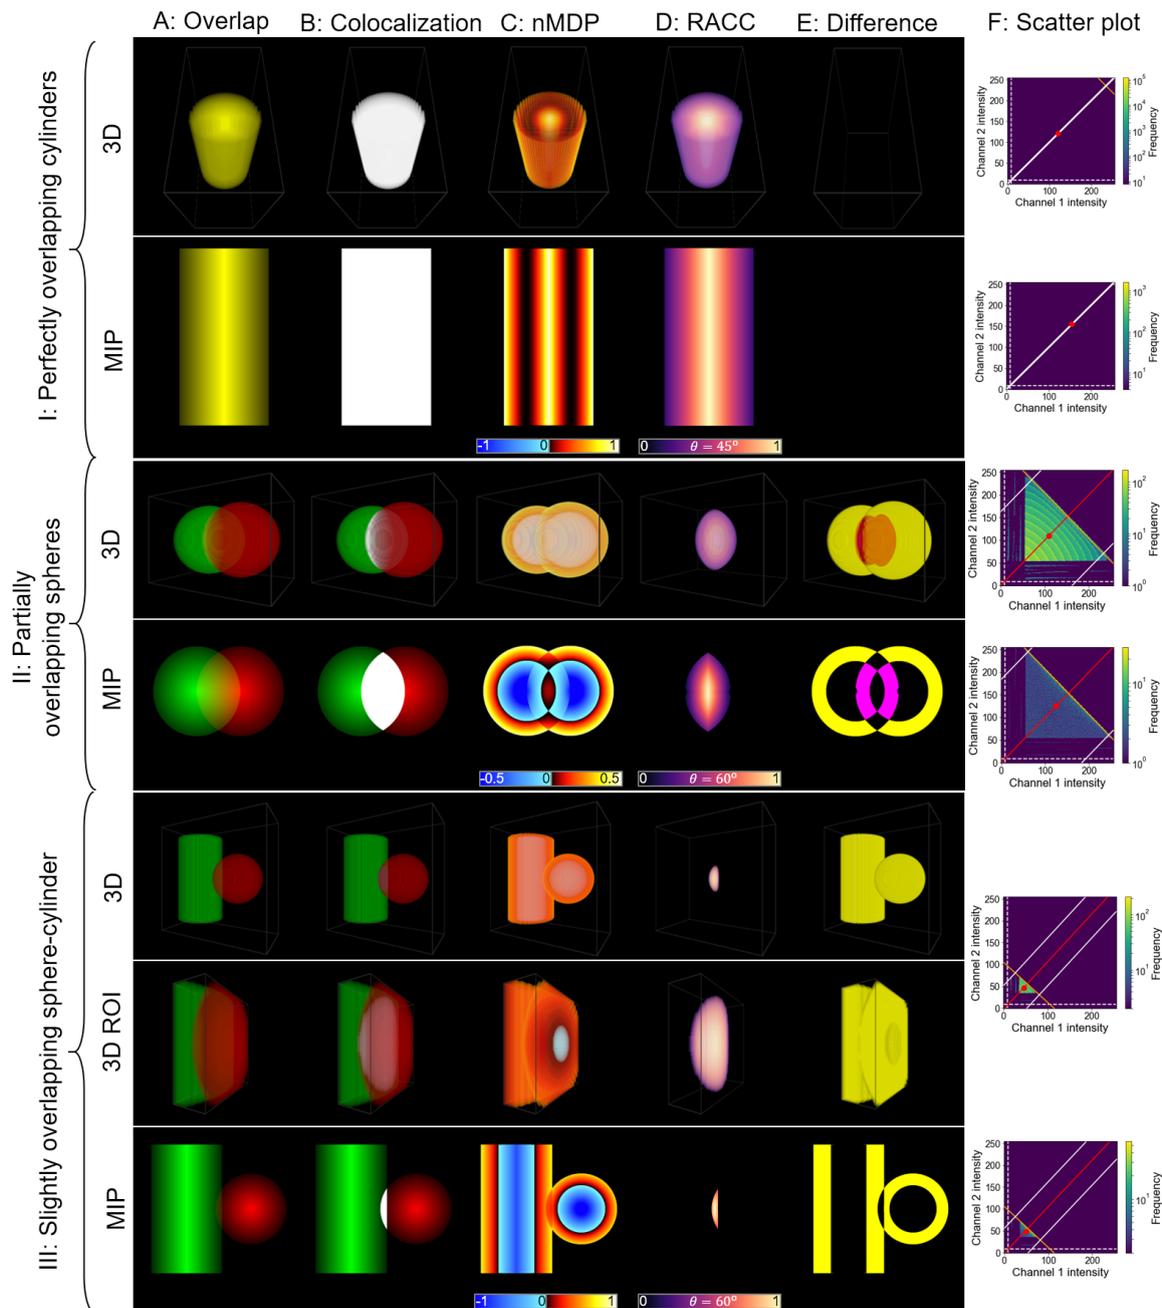


Figure 5.5: The visualisation of synthetic data using both 3D images and maximum intensity projections (MIP). A: The overlapping fluorescence channel intensities. B: All voxels above the colocalisation thresholds and therefore considered colocalised voxels are overlaid in white. C: The result of applying the nMDP. D: The result of applying RACC. E: The difference between RACC and the nMDP. Magenta represents areas which nMDP considers not-colocalised but that RACC considers colocalised, while yellow represents areas that nMDP considers colocalised but RACC does not. F: Scatter plots for the 3D and MIP data sets, showing the frequency of each colocalisation intensity. The regression line calculated by RACC is shown in red. The maximum and distance thresholds for RACC are shown in orange and white, respectively. The red marker represents the per-channel average intensities used by the nMDP calculations and around which the four quadrants are separated (see Figure 5.1).

fluorescence intensity, which typically correspond to a higher concentration of the fluorescent label. Note that since the x and y of all colocalisation intensities are the same and form a line on the scatter plot in Figure 5.5F, the penalisation factor θ has no effect on the visualisation.

When considering the partially overlapping spheres (Figure 5.5 section II), the fluorescence intensities at the centre of the overlapping volume are strongly correlated, with correlation decreasing towards the edges of the overlapping region. This can be easily discerned from the RACC visualisation in column F. Notice also that the left and right sides of the RACC visualisation are slightly clipped due to the distance threshold (Equation 5.18) which removes outliers. Furthermore, since the maximum point is calculated to include 99% of the voxels (indicated by the orange line in Figure 5.5F), the entire colourmap spectrum is used for the visualisation. On the other hand, the nMDP visualisation is more ambiguous, since the region near the surface of the spheres appears to contain the majority of the colocalised voxels, while in fact there is no colocalisation in that region as indicated by column B. Furthermore, the nMDP does not correctly show the entire overlapping volume as colocalised, with the overlapping volume further from the centre being shown as not-colocalised.

When considering the slightly overlapping sphere and cylinder (Figure 5.5 section III), there is only a small region in which colocalisation occurs. Since the intensities of the two channels are similar in the colocalised region, they should be fairly well correlated. This aspect is intuitively reflected when using RACC, where regions of similar fluorescence intensity in the centre of the overlapping region are highlighted. As in the case of the partially overlapping spheres, regions of lower channel intensity near the surface are again erroneously labelled as colocalised by the nMDP. Moreover, the nMDP does not visualise the overlapping volume noticeably different from the surrounding structure. This leads to a visualisation that could be difficult to interpret.

5.3.2 Visualising biological data

The nMDP has been and continues to be used in many studies in the life sciences, where the precise spatial analysis of colocalisation is of particular relevance. These include the assessment of molecular interactions at the neuromuscular junction [88], the analysis of vesicular structures part of the endosomal compartment [97, 98], the characterisation of filamentous actin or tubulin network structures [89, 91] and the localisation and assessment of the degree of protein interaction at the sub-cellular level [90, 92].

In most of these studies, the quantification of colocalisation based on a colourmap is limited and conclusions are based primarily on the comparison of colocalisation metrics such as the PCC. RACC may offer new perspectives, by restricting the analysis to the colocalised voxels and by allowing a more robust identification of the intensity *correlation* at colocalised voxels.

In the following, we apply both nMDP and RACC to a subset of biological samples with similar complexity and compare the results.

5.3.2.1 Cell culture and transfections

Mouse embryonic fibroblasts (MEFs) as well as stably expressing GFP-LC3 MEF's were supplemented with Dulbecco's Modified Eagles Medium (DMEM), 1% penicillin/streptomycin (PenStrep) (Life Technologies, 41965062 and 15140122) and 10% foetal bovine serum (FBS) (Scientific Group, BC/50615-HI) and incubated using a humidified incu-

bator (SL SHEL LAB CO₂ Humidified Incubator) in the presence of 5% CO₂ at 37°C. Where applicable, cells were micropatterned, as previously shown [99]. Cells were seeded in either an 8-chamber cover slip-based dish (Nunc, Lab-Tek, 155411) or a micropatterned slide for subsequent experiments.

5.3.2.2 Immunofluorescence and super resolution structured illumination microscopy (SR-SIM)

In brief, MEF cells were fixed in 4% formaldehyde (Sigma-Aldrich, USA), washed with phosphate buffered saline (PBS) and permeabilised using 0.2% Triton-x, followed by a blocking step using 1% bovine serum albumin (BSA). Thereafter, cells were incubated at 4°C overnight in primary antibodies against α/β tubulin (Cell Signaling, #2148), acetylated tubulin (SC-23950, Santa Cruz) and LysoTracker red (Life Technologies, #L-7528), followed by staining with secondary antibodies Alexa-488 and Alexa-561. SR-SIM and confocal micrographs were acquired using the ELYRA PS.1 station (Carl Zeiss Microimaging; Germany). Thin (0.1 μm) Z-stacks of high-resolution image frames were collected in 3 rotations, followed by reconstruction using ZEN software (black edition, 2011, version 7.04.287) based on a structured illumination algorithm [70].

5.3.2.3 Analysis of α/β tubulin and acetylated tubulin overlap

The tubulin network in a cell is characterised by dynamic modifications such as acetylation, which impacts its properties such as stiffness and function. Here, we assess the degree of acetylation relative to the entire tubulin network through colocalisation analysis. Since microtubulin are delicate filamentous structures, their overlap is very challenging to analyse.

In order to analyse colocalisation, we make specific ROI selections at the periphery of the cell (row II in Figure 5.6) as well as a cross section along the z-axis in the perinuclear region (row III in Figure 5.6). This cross section was introduced so that the colocalisation inside the tubulin can also be visualised, where the correlation between the fluorescence channels is the strongest. In this way, we are able to analyse colocalisation with precision in regions with both high and low fluorescence channel intensity. The location of these selections is shown on the MIP in row IV of Figure 5.6.

In row II of the nMDP visualisation in Figure 5.6C the central area is clearly highlighted as being colocalised. However, since a colourmap value is assigned to all voxels for which a fluorescence channel is present, it is not possible to clearly discern the extent of this colocalisation. RACC similarly highlights the central region as colocalised (row II in Figure 5.6D). However, it also indicates that there is a greater extent of colocalisation in this region. Regions of lower intensity and with a weaker correlation between fluorescence channels are retained when using RACC and are shown as colocalised in darker hues of blue, whereas nMDP reports these regions as not colocalised. Furthermore, nMDP falsely demarcates several regions around edges of the sample as being colocalised, indicated by yellow in Figure 5.6E, which leads to a visualisation that detracts from the truly colocalised regions.

A prominent tubulin strand was isolated from the sample in row III of Figure 5.6D. Here the nMDP indicates no colocalisation, while RACC is able to indicate this volume as colocalised and allows one to discern the varying degree of correlation.

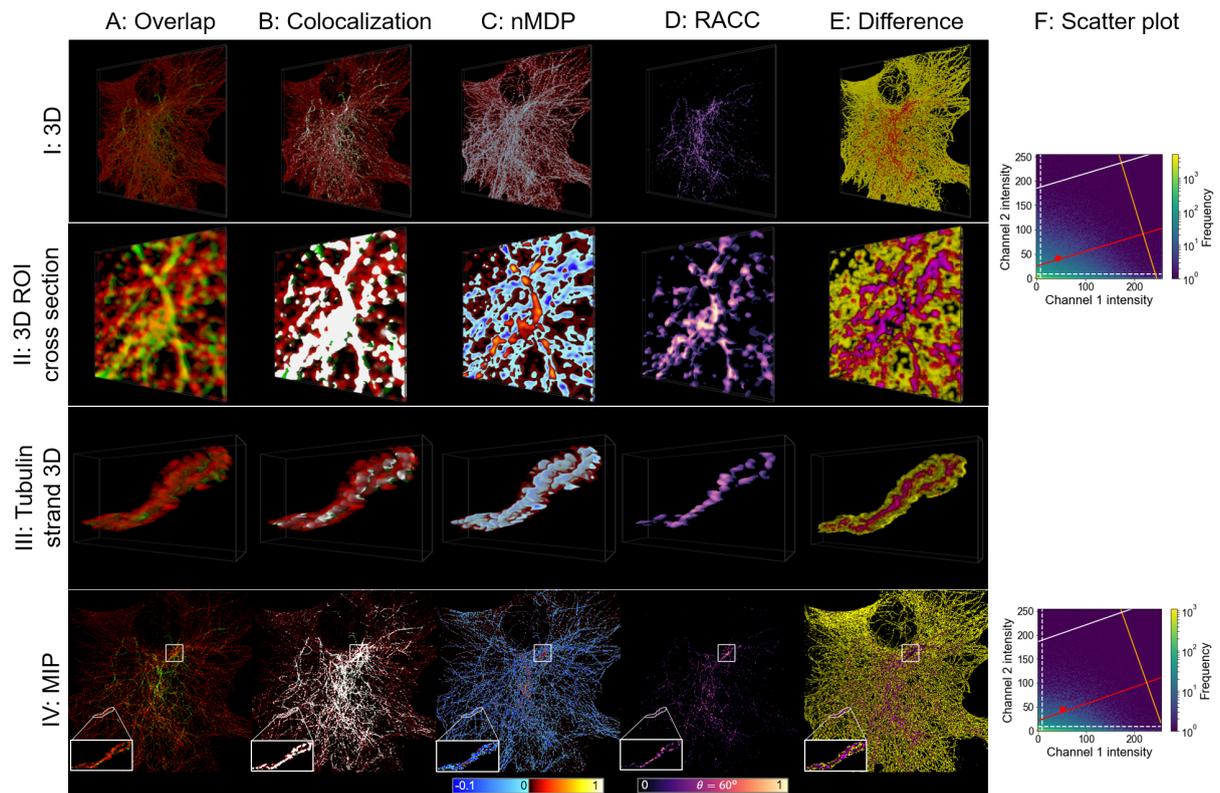


Figure 5.6: The visualisation of α/β tubulin (red) and acetylated tubulin (green). The structure of the figure is similar to that of Figure 5.5, but in this case biological samples and not synthetic data are considered.

5.3.2.4 Analysis of autophagosome and lysosome fusion

Protein degradation through macroautophagy plays a critical role in cellular homeostasis, metabolism and disease [100]. It is characterised by the delivery of autophagosomes to lysosomes, where, upon fusion, an autolysosome is formed, and hydrolytic degradation takes place. To better understand autophagy function and dysfunction, it is of major interest to discern the contribution of the pathway intermediates to the total intracellular vesicle pool size, i.e. autophagosomes, autolysosomes and lysosomes [99]. Visualising the fusion zone between autophagosomes and lysosomes is hence of critical importance, and will be used here as a second example to compare RACC and nMDP (Figure 5.7).

Due to the small size of these organelles, it is challenging to visualise the degree of fusion that has taken place. Note that this fusion is what has been modelled in the synthetic data by means of partially overlapping spheres. We recall from the synthetic data that the nMDP erroneously indicated regions near the surface of the spheres as colocalised. A similar observation is made in the case of the here described autophagosome/lysosome fusion in Figure 5.7C. This may detract from the true interaction between the organelles and hence from the correct interpretation of the colocalisation data. This is also reflected in the substantial difference between nMDP and RACC shown in Figure 5.7E.

RACC shows only the colocalised voxels, which allows the colocalised regions with a strong correlation of the fluorescence channels (high intensities of the RACC colourmap) as well as regions with lower correlation (darker regions) to be clearly discerned. This is best seen in rows II and III of Figure 5.7D. By isolating a small region of the cell, shown in

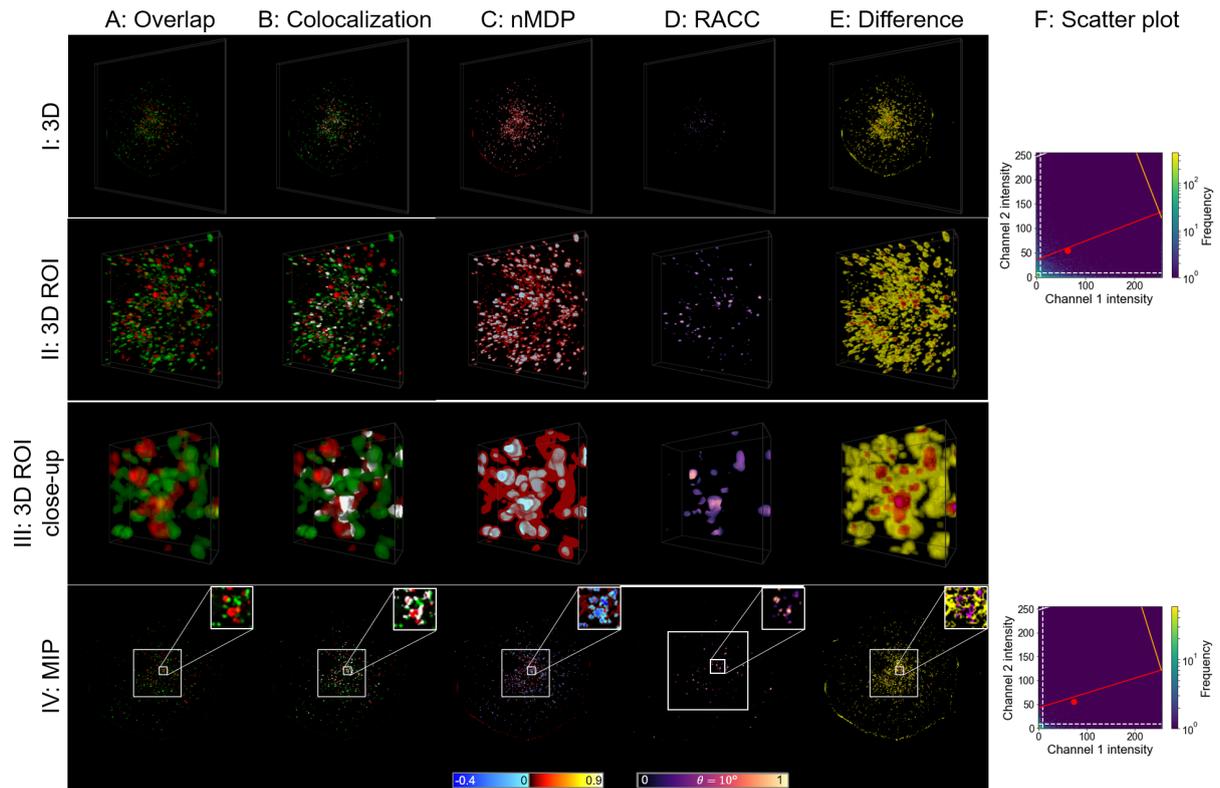


Figure 5.7: The visualisation of autophagosome (green) and lysosome (red) fusion. The structure of the figure is similar to that of Figure 5.5, but in this case biological samples and not synthetic data are considered.

row III, RACC reveals which of these colocalisations indicates an almost complete fusion between the organelles and hence the extent of autophagy progression.

5.3.2.5 Analysis of autophagosome and tubulin interaction

Autophagosomal transport is facilitated by the microtubulin network [101]. Hence, the interaction between these two structures and its accurate visualisation is of major interest.

In Figure 5.8 we show such autophagosome and tubulin interaction. Since the autophagosomes are transported along the tubulin network, there is very little overlap between them and consequently very few colocalised voxels. We therefore assess two ROIs in order to investigate the interaction in more detail. These ROIs are shown in rows II and III of Figure 5.8 and their location within the sample as a whole is shown on the MIP in row IV.

As for autophagosomes and lysosomes (Figure 5.7C), the nMDP in Figure 5.8C does not clearly indicate the correlation of the colocalisation between the fluorescence channels. This is mainly due to the incorrect identification of voxels at the edge of the structures as colocalised. A similar result was also observed for the synthetic data consisting of partially overlapping spheres.

The magenta areas in row III of Figure 5.8E, which appear red due to the yellow surrounding voxels, also match the structures that RACC identifies as colocalised in Figure 5.8D. This indicates that nMDP is not suitable to precisely and correctly visualise the truly colocalised regions, shown in white in Figure 5.8B. This can be ascribed to the

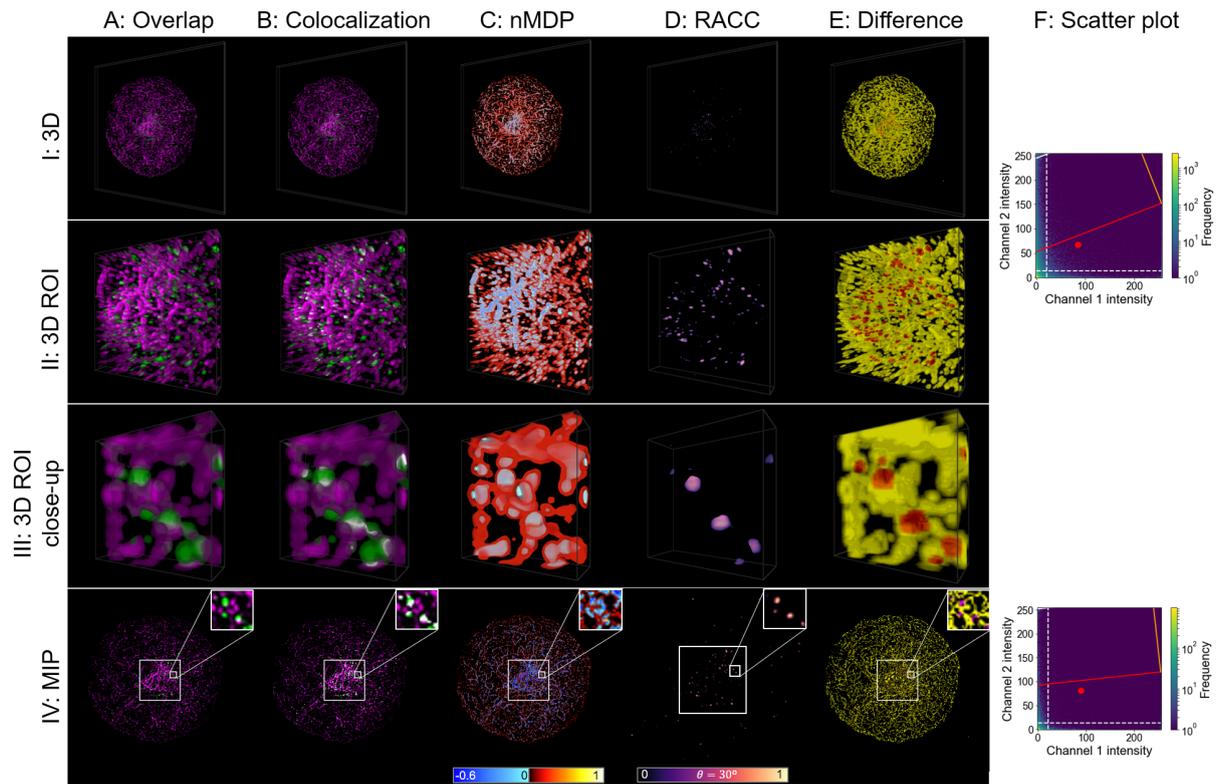


Figure 5.8: The visualisation of autophagosome (green) and tubulin (magenta) interaction. The structure of the figure is similar to that of Figure 5.5, but in this case biological samples and not synthetic data are considered.

lack of correlation between the fluorescence channels, as can be seen in the scatter plots. As a result, most of the high intensity voxels, located inside the structure, are labelled as not-colocalised (blue). The lower intensity voxels, which surround the structure, are labelled as colocalised (red). However, there is little true colocalisation in this sample, and the weak true colocalised voxels localises in the higher-intensity region. nMDP is not able to adequately compensate for this imbalance, contributing to the observed differences. The percentage of the total voxels that were classified differently by RACC and nMDP is shown in Table 5.1.

Since RACC shows only colocalised voxels, by considering it in conjunction with the overlapping fluorescence channels (Figure 5.8A and B), a better understanding of the interaction between autophagosomes and tubulin can be obtained. The qualitative extent of these colocalisation events can also be determined more intuitively.

5.3.3 Discussion

Although methods to visualise the correlation between two colocalised fluorescence channels in a qualitative way exist, challenges remain when interpreting these visualisations because of ambiguity introduced into the visualisation due to not properly considering the colocalisation of the fluorescence channels. Here we have demonstrated how RACC can offer an improved interpretability of the nature of the colocalisation in different regions of a sample. The synthetic and biological sample data used serve to indicate that RACC can be applied to a variety of colocalisation investigations, including samples that

are challenging to analyse. As the scatter plots that accompany the figures show, RACC enables the visualisation of data even when no clear correlation between the fluorescence channel intensities can be discerned.

Even though RACC can be used as a stand-alone analysis tool, it may be more powerful when utilised in conjunction with the popular overlapping fluorescence channel visualisation (sub-figures A and B in Figures 5.5-5.8), since these methods contextualises the colocalisation within the surrounding cell structure.

The RACC parameters that were used when visualising the samples presented in this chapter were calculated over the sample as a whole. Since the ROIs that we isolated were only intended to show the differences between nMDP and RACC in greater detail and not to make biological conclusions, we chose not to recalculate the RACC parameters for individual ROIs. However, such recalculations based on careful ROI selections are strongly advised when region-specific analysis is required, since each sub-cellular region can have different colocalisation correlation distributions, and consequently different regression lines. Furthermore, care should be taken to minimise additional factors that could influence the fluorescence channel intensities, such as the relative quantum yield of fluorophores, the collection efficiency of the microscope and the detector gain [60].

An implementation of RACC, allowing the visualisation of both 2D and 3D samples, is available for download at <https://rensutheart.github.io/RACC.html>.

5.4 Conclusion

Visual colocalisation analysis in fluorescence-based microscopy is a very important method used by biologists to gain an understanding of the association between molecular structures within a cell. However, existing methods of colocalisation visualisation that qualitatively visualise the spatial distribution of the correlation between the fluorescence channels in a sample have several limitations which could lead them to produce ambiguous visualisations in certain situations, thereby reducing their usefulness.

To overcome these limitations, we present RACC (regression adjusted colocalisation colour mapping), a novel biological visual analysis method that offers improved spatial visualisation of colocalisation in fluorescence-based micrographs using a colourmap. RACC is designed to meet three objectives: firstly to highlight regions within a cell that have a strong positive correlation between two fluorescence channels, secondly to highlight colocalised regions that have greater combined fluorescence intensities and thirdly to suppress coincidentally colocalised voxels from the visualisation. These objectives were achieved by applying Deming regression in conjunction with a geometric projection of the fluorescence intensities while using automatically calculated thresholds to remove outliers.

By visualising both synthetic data and three biological samples, RACC was compared with a prominent existing method used to visualise the fluorescence channel intensity correlation of colocalised volumes using a colourmap. These visualisations demonstrated how RACC can offer an improved understanding of the nature of the colocalisation within a sample by visualising only truly colocalised regions while offering a qualitative measure of the correlation between fluorescence channels in different regions of a sample. This positions RACC as an attractive colocalisation analysis tool especially in scenarios characterised by complex structure and biology.

It is hoped that the application of RACC in fluorescence microscopy will enable better discrimination of colocalisation events and in turn assist future biological investigations where colocalisation analysis is central.

In the next chapter we continue to extend our virtual reality based system to incorporate the ability to analyse fluorescence microscopy based time-lapse image sequences. We specifically investigate the mitochondrial fission, fusion and depolarisation events by developing a new method to automatically localise and quantify these events.

Chapter 6

Mitochondrial event localiser (MEL) to quantitatively describe fission, fusion and depolarisation in 3D

The content of this chapter was submitted to PLOS one as a research article:

Theart, R.P., Kriel, J., Loos, B. and Niesler, T.R., 2019. Mitochondrial event localiser (MEL) to quantitatively describe fission, fusion and depolarisation in the three-dimensional space.

6.1 Introduction

Mitochondria are highly dynamic organelles that do not operate in a stagnant or isolated manner. Rather, they function in a highly energetic and networked fashion, continuously subjected to rapid remodelling events referred to as fission (fragmentation) and fusion [102]. These events serve as a critical quality control mechanism, not only to adapt and respond to changing metabolic demands but also to enable separation of damaged mitochondria from those in an interconnected state and thereby exposing them to specific degradation [103]. The elimination of damaged mitochondria, mediated through a mitochondrial specific autophagy referred to as mitophagy, decreases the risk of mitochondrial DNA (mtDNA) mutation accumulation. It also improves the electron transport chain (ETC) efficiency [104, 105, 106, 107]. The transport of electrons between different complexes on the mitochondrial membrane results in the maintenance of a continuous voltage across the inner mitochondrial membrane (IMM) referred to as the mitochondrial membrane potential ($\Delta\psi_m$) [108]. Mitochondrial dynamics are highly dependent on this membrane potential. Excessive membrane potential dissipation has, for example, been observed to render fragmented mitochondria incapable of re-fusing, thereby metabolically debilitating the cell [108, 109, 110]. Furthermore, the recruitment of the mitophagy initiation protein PTEN-induced kinase 1 (PINK1) only occurs once a certain depolarisation threshold is reached [111]. Failure to eliminate damaged mitochondria is a hallmark of neurodegenerative diseases such as Parkinson's disease [112], while certain heritable diseases, including Charcot-Marie-Tooth Type IIA, are related to dysregulated mitochondrial dynamics [113]. This underscores the important role mitochondrial fission and fusion play, not only in maintaining mitochondrial homeostasis, but also in preserving overall cellular

viability [114]. It is thus not surprising that there is a lot of interest among researchers to quantify these dynamic changes accurately.

Previous attempts to better describe and unravel the interplay between mitochondrial dynamics and cell death onset have largely focused on either the change in the mitochondrial network of the entire cell or the rate at which fission and fusion occurs [112, 115, 116]. The accurate description of the quantitative relationship between fission and fusion dynamics in a three-dimensional cellular context, in order to detect deviation from its equilibrium, has however remained challenging. Moreover, although it is becoming increasingly clear that intracellular localisation of mitochondria is indicative of regional specific functions, for example the reliance of cellular organelles on mitochondrial ATP provision, it remains largely uncertain where in the mitochondrial network depolarisation is most likely to occur. It therefore remains to be determined which areas of the mitochondrial network are preferentially depolarised to facilitate either transportation or degradation and how these areas relate to the mitochondrial morphometric parameters usually employed. Currently, the methodologies available to address this question are limited, particularly in the context of three-dimensionally-based quantification of fission and fusion dynamics concomitantly with the onset of mitochondrial depolarisation. To gain a better understanding of mitochondrial dynamics, a high-throughput, automated and deterministic method that is able to localise and quantify the number of mitochondrial events in large three-dimensional (3D) time-lapse sample sets would therefore be advantageous. Such a method would allow subsequent quantitative description of the fission/fusion equilibrium as well as the extent of depolarisation.

Using time-lapse microscopy data of cells labelled for mitochondria, it has previously been observed that fission and fusion are rapid events occurring within a five second time frame under homeostatic conditions [117]. This work, however, did not investigate mitochondrial depolarisation. Current methods typically rely on the manual comparison of two time-lapse image frames in order to observe where fission, fusion or depolarisation occurred. This very labour intensive approach makes it challenging to gain comprehensive insights into the mitochondrial dynamics of a whole three-dimensional sample. It is usually unclear whether mitochondrial fission and fusion events are changing spatio-temporally, thereby shifting the equilibrium towards either fission or fusion. In this context, it is also often not clear whether a cell with a more extensively fused mitochondrial network is in transition or whether it has established a new equilibrium between fission and fusion by means of a newly adapted net contribution of fusion events. Similarly, a cell with greater mitochondrial fragmentation may establish a new equilibrium between fission and fusion as an adaptive response, in order to maintain and preserve this fragmented morphology. Hence, the micrograph reflects an underlying relationship between fission and fusion and this relationship is currently challenging to describe. Although mitochondrial photoactivation [115, 116] provides a highly selective tool to quantitatively assess mitochondrial dynamics, it does not reveal the relative contribution of fission and fusion events to the observed dynamics. For example, it has been noted that mitochondrial fusion is often enhanced during adaptations to metabolic perturbations, particularly in the perinuclear region, and may protect mitochondria from degradation [118]. On the other hand, a degree of fragmentation is required and desirable to allow, for example, mitochondrial transport in neurons to reach synaptic connections. Yet, describing the precise occurrence of fission, fusion and depolarisation has remained challenging. Although quantitative analysis to understand the dynamic nature of mitochondrial movements and events and fission and fusion better has received significant attention in recent work [117, 119, 120, 121, 122, 123, 124], there is, to our knowledge, currently no method-

ology available that automatically localises mitochondrial events in a time-lapse image sequence.

To address these current challenges, we have developed an approach that automatically localises and count the number of mitochondrial events occurring between two micrograph frames in a time-lapse sequence. The results of this automatic analysis allow the quantitative assessment of fission, fusion and depolarisation localisation with high accuracy and precision. We refer to this newly developed procedure as the *mitochondrial event localiser* (MEL), because it allows us to indicate the precise three-dimensional location at which fission, fusion and depolarisation are likely to occur next. Moreover, MEL enables the determination of the individual locations of smaller structures that fuse to form a larger central structure in the two time-lapse frames that are considered. Similarly, the locations of smaller structures that will separate from a common central structure due to fission are identified. In doing so, MEL provides a platform to better understand mitochondrial dynamics in the context of health and disease, with both screening and diagnostics potential. This approach is, as far as we are aware, the first automated method for the detection of depolarised mitochondria in the context of fission and fusion events. MEL can serve both as a standalone method or as part of the broader mitochondrial analysis pipeline to enable high-throughput analysis of time-lapse data.

6.2 Method

The development of MEL was inspired by previous work which applies a “vote casting” methodology to individual mitochondrial structures in two consecutive time-lapse frames in order to identify the likely mitochondrial event (fission, fusion, or none) the structure will undergo [117]. The purpose of this approach was specifically to count the number of fission and fusion events and can not easily be extended to determine their location. MEL was designed from the outset to both count the number of mitochondrial events and to localise each in 3D space.

The MEL algorithm processes a fluorescence microscopy time-lapse sequence of z-stack images that were stained for mitochondria and produces the 3D locations of the mitochondrial events occurring at each time step. These locations can subsequently be superimposed on the z-stacks in order to indicate the different mitochondrial events. The algorithm has been organised into two consecutive steps, namely the image pre-processing step which normalises and prepares the time-lapse frames, and the automatic image analysis step which calculates the location of the mitochondrial events based on the normalised frames.

6.2.1 Image pre-processing

The image pre-processing step receives a time-lapse sequence of z-stacks as input and begins by selecting two z-stacks, which we will refer to as Frame 1 and Frame 2, for further processing. Depending on the temporal resolution that is desired, the selected z-stacks can either be consecutive time-lapse frames, or some number k of intermediate frames may be skipped. Since mitochondrial movement is highly dynamic, a time interval that is too long might prevent the algorithm from accurately matching mitochondrial structures between the chosen frames, thereby degrading the ability to determine mitochondrial events. The selected Frames 1 and 2 are then each processed in the same way by the image pre-processing step to generate several new image stacks which are passed to the

automatic image analysis step. The image pre-processing step, which will be described in the remainder of this section, is illustrated in Figure 6.1.

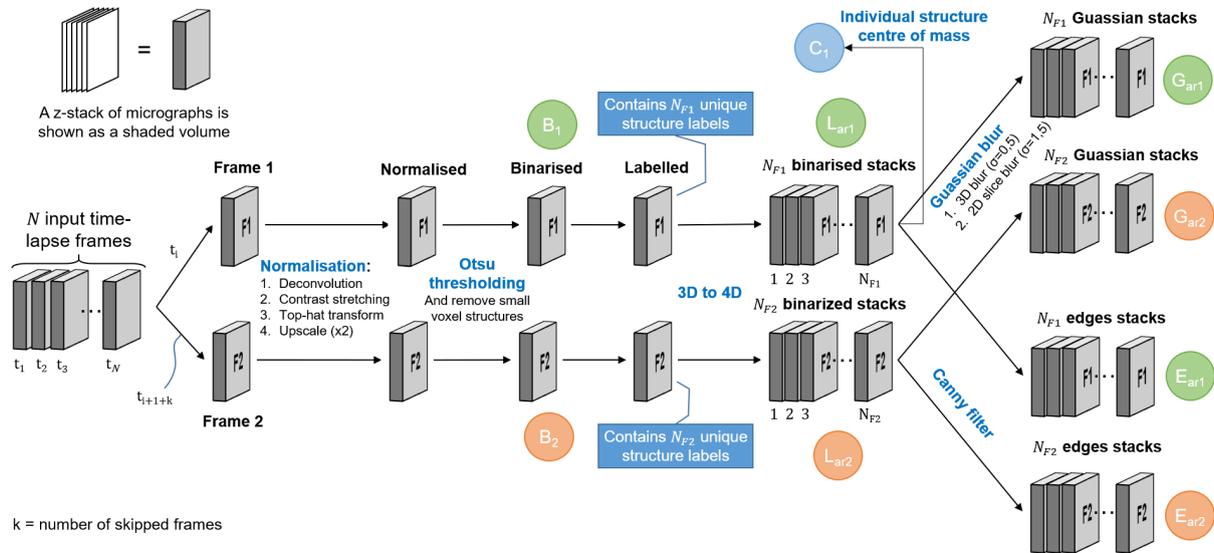


Figure 6.1: Image pre-processing begins by choosing two frames from the input time-lapse sequence, Frame 1 and Frame 2. Normalisation, binarisation, and labelling are then performed on both frames. The labelled frames are each separated into an array of z-stacks, where each stack contains only a single labelled structure. Each of the stacks in the array of binarised stacks, for both Frames 1 and 2, is then Gaussian blurred and Canny filtered. The latter contains only the edges of the labelled structures. The stacks that are used during the automatic image analysis step are labelled in green for Frame 1 and orange for Frame 2. Finally, an array containing the centre of mass for each labelled structure is shown in blue.

Since MEL is not based on the analysis of a single cell, it is not necessary to select regions of interests (ROIs) before the analysis. However, since the image acquisition parameters that are used vary widely between different time-lapse sequences, we first normalise the fluorescence intensity data of the z-stacks. This is necessary since MEL relies on an accurate binarisation of the fluorescence image stack to identify the voxels that contain mitochondria. These thresholding algorithms require images to be sharp, contain minimal noise, and have good contrast between the foreground and background. The normalisation and binarisation, which form the initial part of the MEL image pre-processing step, were guided by previous work focusing on quantitative mitochondrial network analysis [117, 122, 123, 124].

The first step in normalising the raw micrographs is to apply deconvolution to the z-stacks using a point spread function (PSF) that is estimated from the microscope's acquisition parameters. This step removes the out-of-focus light in microscopy images that was induced by the limited aperture of the objective, thereby improving the signal-to-noise-ratio (SNR). The PSF is a model of the 3D impulse response of the fluorescence microscope system, as well as the sample acquisition parameters, and can be obtained either experimentally, analytically, or computationally [125]. For this, we used Huygens Professional deconvolution software. Next, we apply contrast stretching to the z-stack to normalise the fluorescence intensity between micrographs. The micrographs constituting the z-stack are then upsampled by a factor of two using bilinear interpolation so as to

increase the resolution of the binarised image and reduce the possibility that two adjacent but unconnected structures are erroneously joined after binarisation. Finally, we apply a three-dimensional top-hat transformation which further reduces noise and enhances the mitochondria in such a way that they can be more easily isolated from the background.

The normalised frames are binarised by applying Otsu thresholding [126] to the z-stack, although other thresholding methods might also be suitable. Since it has been observed that some noise remnants were also being binarised, we removed any structure containing less than some appropriately small number of voxels (in our case 20) in the upscaled images. This z-stack that results from this is labelled B in Figure 6.1. Note that we differentiate between the stacks of Frame 1 and Frame 2 in Figure 6.1 by using subscripts 1 and 2, respectively.

After binarisation, we assign a label to each separate 3D voxel structure. The total number of labels in Frame 1 and Frame 2, as depicted in Figure 6.1, are N_{F1} and N_{F2} , respectively. Later in the automatic image analysis step, each labelled structure in each frame will be compared with labelled structures both within the same frame as well as in the other frame. For this reason, we separate the labelled z-stack into an array of z-stacks, where the i^{th} stack in the array contains only the binarised voxel structure associated with label number i . We refer to this array of stacks as L_{ar} . The subscript “ ar ” is used to indicate an *array* of stacks, containing N_{F1} and N_{F2} stacks for Frame 1 and Frame 2, respectively. This step can be thought of as adding the structure label as a fourth dimension to the data and later enables fast comparison between different labelled structures.

From the array of labelled z-stacks L_{ar} we create two similar arrays also required by the automatic image analysis step. The first is the result of applying a 3D Gaussian filter to each z-stack in the array. This slightly blurs the edges of the labelled stacks and results in the array of z-stacks G_{ar} . The Gaussian blurring enhances the ability of the algorithm to match the moving mitochondrial structures between two frames by slightly inflating the structures. The second array of stacks is generated by removing all voxels that are not located on the edges of the 3D labelled structures in L_{ar} by using Canny edge detection [127]. We refer to the resulting array of z-stacks as E_{ar} . This is performed mainly to improve the efficiency of the algorithm that later determines the location of fission and fusion events.

This concludes the image pre-processing step. The tuneable parameters in this step is the volume of the structures that are considered as noise, and the standard deviation with which the z-stacks are Gaussian blurred. The z-stack B , and the arrays of z-stacks G_{ar} and E_{ar} are now passed into the automatic image analysis step that calculates the location of the mitochondrial events.

6.2.2 Automatic image analysis step

The purpose of the automatic image analysis step is to generate a list of locations in Frame 1 at which the mitochondrial events occur. It achieves this by receiving the z-stacks that were generated in the image pre-processing step. These are then automatically analysed to produce a list of the mitochondrial event locations. The automatic image analysis step is shown in Figure 6.2.

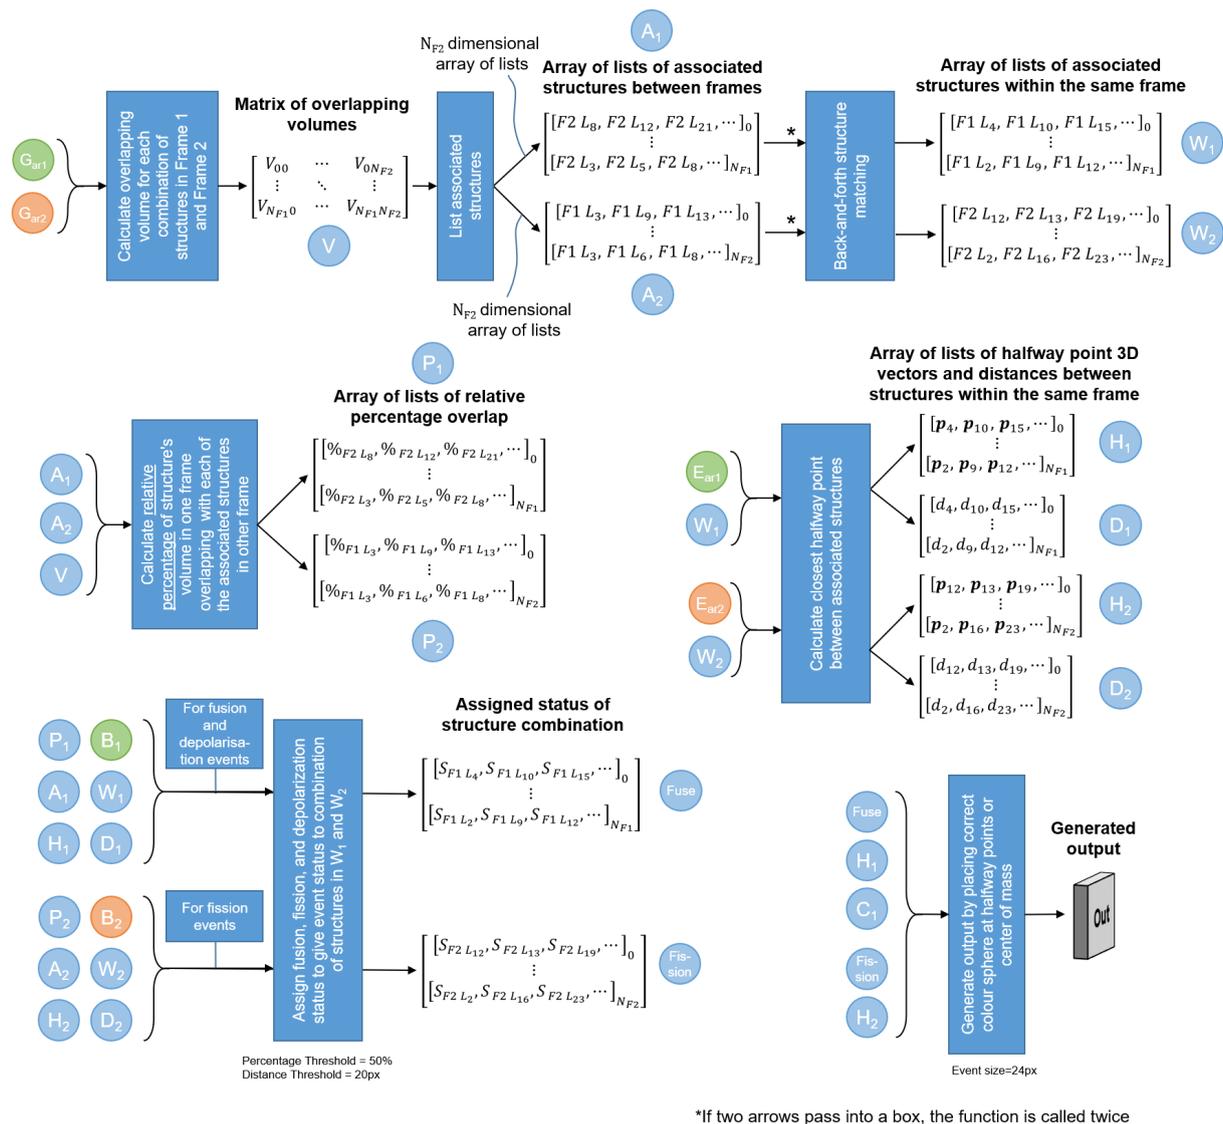


Figure 6.2: Automatic image analysis begins by calculating the overlapping volume of each structure in Frame 1 with each structure in Frame 2. V_{xy} refers to the number of voxels that overlap between structure x in Frame 1 and structure y in Frame 2. To compensate for such coincidental matches, all overlapping volumes that account for less than 1% of the volume of either structure in question are eliminated. Label 0 refers to the background. Using matrix V , the labelled structures that are associated with each other both between Frames 1 and 2, as well as within the same frame, can be calculated. The relationships are encoded as arrays of lists, where each list is related to the labelled structure corresponding to the array index. Therefore, the 8th list in the array relates to labelled structure number 8. $F2L_8$ refers to label number 8 in Frame 2 that is associated with the structure at the array index. Next, the relative percentage of overlap (e.g. $\%_{F2L_8}$) is calculated for the associated structures between the frames, and is interpreted in conjunction with A_1 and A_2 . The midway points (e.g. p_4) and distances (e.g. d_4) between two structures within the same frame are calculated, and are interpreted in conjunction with W_1 and W_2 . Using these arrays of lists, the status (e.g. S_{F1L_4}) of fission, fusion, depolarisation or no event can be determined for each mitochondrial structure. Finally, the event can be overlaid on the original z-stack at the location of the midway points.

6.2.2.1 Mitochondrial events

The three different mitochondrial events that we consider, namely fission, fusion and depolarisation, are defined in terms of changes in mitochondrial morphology that occur in the time interval between Frames 1 and 2. Specifically, fission is defined as the separation of a larger mitochondrial structure to two smaller structures. Similarly, fusion is defined as the joining of two mitochondrial structures to form a single larger structure. Finally, depolarisation is defined as the disappearance of a mitochondrial structure from Frame 1 to Frame 2 manifested by a complete loss of fluorescence signal. If more than two structures fuse to a single central structure, or if one structure undergoes fission to form more than two smaller structures, then separate locations are assigned to each.

6.2.2.2 Process

In order to determine the location of mitochondrial events, we first identify which of the voxel structures labelled during the image pre-processing step are candidates for fission and fusion. First, the voxel structures in Frame 1 that occupy the same space as a voxel structures in Frame 2, and vice versa, are identified. For the sake of conciseness, we will refer to such collocation of structures between the two frames as overlap. Using this overlap information, we can determine which voxel structures are associated with each other within the same frame. In Frame 1, these associated structures will be the fusion candidates, and in Frame 2 they will be the fission candidates. Structures in Frame 1 that have no associated structures in Frame 2 undergo depolarisation.

To determine the associated structures between Frame 1 and Frame 2, we first calculate a matrix V containing the overlapping volume (in voxels) of each structure in Frame 1 with each structure in Frame 2 using the blurred arrays of z-stacks G_{ar_1} and G_{ar_2} . Label number 0 is reserved for the background. Blurred stacks are used in order to allow for movement of the mitochondria between frames. Using matrix V , we calculate two arrays A_1 and A_2 . Each entry in array A_1 is a list of all the structures in Frame 2 that overlap with a particular structure in Frame 1. Similarly, A_2 indicates the structures in Frame 1 that overlap with a particular structure in Frame 2.

The information in A_1 and A_2 allow us to determine which structures in Frame 1 are fusion candidates, and which structures in Frame 2 are fission candidates using an algorithm we call *back-and-forth structure matching*. For each structure in Frame 1, A_1 is used to determine all associated structures in Frame 2. For each of these associated structures in Frame 2, A_2 is used to determine all the associated structures in Frame 1. This list of Frame 1 structures comprises the fusion candidates for the Frame 1 structure being considered. An identical but opposite procedure is used to determine the fission candidates in Frame 2. The resulting two arrays of lists are denoted W_1 and W_2 . In the remainder of automatic image analysis step we will determine which of these candidates truly underwent the mitochondrial event. The back-and-forth structure matching algorithm is illustrated in Figure 6.3 and is practically demonstrated with the help of a synthetic example in Appendix C.

Many of the mitochondrial event candidates identified in the previous step might be a result of coincidental structure overlap. These overlaps could either be due to mitochondrial movement in the time interval between Frame 1 and Frame 2, or as a result of comparing the blurred frames G_{ar_1} and G_{ar_2} . Blurring was introduced to compensate for mitochondrial movement, but has the consequence of increasing the number of false overlaps. To remove these coincidental overlaps, we begin by calculating the relative per-

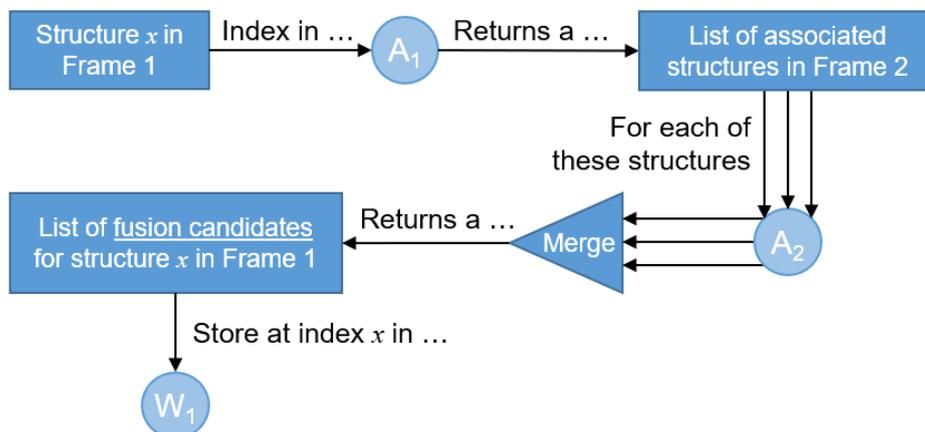
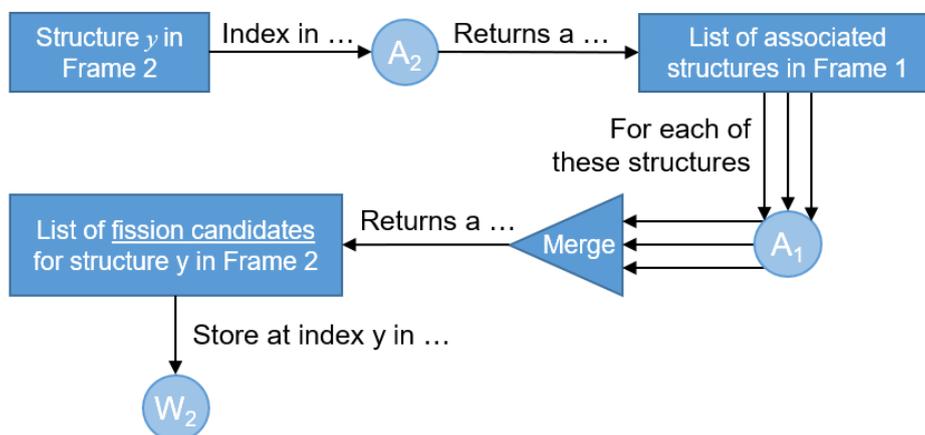
A: Fusion candidate determination**B: Fission candidate determination**

Figure 6.3: The back-and-forth structure matching algorithm used to determine mitochondrial structures that are candidates for fission and fusion.

centage overlap between associated structures in Frame 1 and Frame 2. For each structure in Frame 1, the set of associated structures in Frame 2 is identified using A_1 . Then, using V , the overlapping volume of each of these associated structures is normalised relative to the combined overlapping volume of all the associated structures. This results in the relative percentage overlap that each structure in Frame 2 has with the structure in Frame 1. In this way structures with a small percentage overlap can be considered coincidental and ignored as candidates when determining the type of mitochondrial event. These overlap percentages are stored in an array of lists P_1 , which has the same structure as A_1 . An analogous procedure is followed for all the structures in Frame 2, in this case using A_2 , resulting in P_2 . P_1 and P_2 are used to remove mitochondrial event candidates whose proportion of volume overlap is small.

The back-and-forth structure matching algorithm can also produce false matches. For example, when two small structures are on opposite sides of a larger central structure in Frame 1 and both fuse with this central structure in Frame 2, the smaller structures in Frame 1 will be identified as being associated with each other. This is, however, a false match since no mitochondrial event occurs between them. Such false matches can easily

be avoided by detecting the presence of a third mitochondrial structure between the two candidates for association, or by determining that the candidate structures are too far apart for a mitochondrial event to be feasible. We therefore calculate the shortest vector length as well as the point in 3D space midway between all structure pairs identified in W_1 and W_2 . This midway point represents a likely location of fission and fusion events. The shortest vector lengths can be computed from the information in E_{ar_1} and E_{ar_2} . The resulting list of shortest distances are denoted by D_1 and D_2 and the midway points by M_1 and M_2 for fusion and fission respectively. Note that the length of each array in the lists matches those of W_1 and W_2 for the same subscript.

Using A_1 and A_2 , we can find the structures in Frame 1 that have no associated structures in Frame 2, indicating that they will depolarise. Next, the false fission and fusion candidate structures in W_1 and W_2 can be filtered out by using D_1 , D_2 , M_1 and M_2 to determine the final sets of fission and fusion events. First, a distance threshold corresponding to the maximum distance that a mitochondrial structure could be expected to move is determined. Then, all candidates in W_1 and W_2 for which the respective distances in D_1 and D_2 exceed this threshold are removed. Next, fission and fusion candidates separated by a third structure are removed by considering the binarised stacks B_1 and B_2 in conjunction with the midway points in M_1 and M_2 . Finally, by using P_1 , we consider two structures within Frame 1 that are both associated with the same structure in Frame 2 with a high relative percentage overlap as a fusion event. Similarly, using P_2 , we consider two structures within Frame 2 that are both associated with the same structure in Frame 1 with a high relative percentage overlap as a fission event. Any structure pairs with a low and high, or a low and low, combination of relative percentage overlap is ignored. Empirically we have observed that in most cases the relative percentage overlap is either above 90% or below 10%. Therefore, we consider any percentage above 50% to be high. This step results in two more arrays of lists. The *Fusion* array indicates which structures in Frame 1 will fuse, depolarise, or undergo no event. The *Fission* array indicates which structures in Frame 2 are the result of a fission event.

The final step is to generate an output image by superimposing a colour label, in three-dimensional space, on the z-stack to indicate the location of each mitochondrial event at the midway point between the two structures that fuse in Frame 1, or using the midway point between the two structures that are the result of fission in Frame 2 as the location at which a structure in Frame 1 will undergo fission. The structures that will depolarise are indicated by placing a colour label at its centre of mass, which is calculated from the z-stack for the structure in question from the L_{ar_1} array.

The automatic image analysis step, indicating which stacks and list of arrays are used for which parts of the algorithm, is shown in Figure 6.2. The tuneable parameters in this step is the volume overlap threshold, the percentage that is the threshold between a low and high relative overlap percentages and the distance threshold between structures that can be considered to be related to each other. In order to further clarify each step in the MEL procedure, we have included its application to a simple synthetic sample in Appendix C.

6.2.3 Biological sample investigation

Using a mammalian cell model, we apply and validate MEL by assessing the mitochondrial network under physiological as well as disrupted conditions. For this we analysed two control cells and contrast these with cells treated with hydroxychloroquine sulphate (HCQ), leading to mitochondrial dysfunction [116].

6.2.3.1 Cell Line Maintenance

U-118MG cells were purchased from the American Type Culture Collection (ATCC) and supplemented with Dulbecco's Modified Eagles Medium (DMEM), 1% penicillin/streptomycin (PenStrep) (Life Technologies, 41965062 and 15140122) and 10% foetal bovine serum (FBS) (Scientific Group, BC/50615-HI) and incubated in a humidified incubator (SL SHEL LAB CO₂ Humidified Incubator) in the presence of 5% CO₂ at 37°C.

6.2.3.2 Microscopy

Live cell confocal microscopy of mitochondrial fission and fusion events was conducted using a Carl Zeiss Confocal Elyra S1 microscope with LSM 780 technology. Prior to imaging, U-118MG cells were seeded in Nunc[®] Lab-Tek[®] II 8 chamber dishes and incubated with 800nM tetramethylrhodamine-ethyl ester to allow for the visualisation of the mitochondrial network (TMRE, Sigma Aldrich, 87917) for 30 minutes in the presence of 5% CO₂ at 37°C. In order to perturb the mitochondrial network, cells were exposed to 1mM hydroxychloroquine sulphate (HCQ) (Life Technologies, T669). HCQ has previously been shown to result in fragmentation of the mitochondrial network by disrupting electron transport chain efficiency [116].

For the two control samples the time interval between frames were 6s and 20s respectively. For the HCQ treated samples the time interval was 61.7s.

6.2.3.3 Visualising the confidence of mitochondrial events

The accuracy of the results that are produced with MEL can be affected by image quality. Therefore, to compensate for the variable image quality of the biological samples, we used MEL to calculate the position of all mitochondrial events for three different contrast stretching parameters in the pre-processing phase. This leads to three output images, each with markers indicating the location of all mitochondrial events. The output images were then divided into non-overlapping 10-by-10-by-2 voxel blocks (x, y and z-axes respectively).

Within each such block, the number of times each type of mitochondrial event was detected was counted for the three different pre-processing parameters. We then visualise the different types of mitochondrial events using different colour labels, where the intensity of these labels is proportional to the relative frequency with which each type of mitochondrial event was detected within each 10-by-10-by-2 voxel block. The intensity therefore serves as an indication of how confident MEL is that the mitochondrial event is occurring at that location.

6.3 Results

The ability to automatically locate and count the number of mitochondrial events that occur in a sample is key to the investigation of mitochondrial dynamics. In this section, we assess and validate MEL by applying it to four different time-lapse sequences of mammalian cells that represent a range of example cases.

Figures 6.4 and 6.5 demonstrate the application of MEL to two control time-lapse sequences of cells under physiological conditions and Figures 6.6 and 6.7 consider the application of MEL to two cells that were treated with HCQ. The control cells were acquired with relatively short time intervals between the image frames, allowing individual

Table 6.1: Mitochondrial event percentages for two control samples and two cells treated with HCQ at four different time points corresponding to column A in Figures 6.4-6.7.

	Time	Average total structures	Fusion %	Fission %	Depolarisation %
Control 1	0s	80.67	21.1%	12.8%	3.3%
	23s	83.33	18.8%	11.2%	6.4%
	45s	87.00	25.7%	20.7%	3.1%
	66s	83.00	18.5%	20.1%	10.8%
Control 2	0s	49.00	41.5%	17.0%	8.8%
	80s	45.67	32.8%	27.7%	12.4%
	160s	47.00	21.3%	33.3%	17.0%
	240s	43.00	42.6%	23.3%	10.1%
HCQ 1	0s	94.33	33.9%	26.9%	43.1%
	248s	47.67	46.9%	25.9%	12.6%
	494s	50.00	70.7%	10.7%	36.0%
	742s	32.67	32.7%	12.2%	17.3%
HCQ 2	0s	71.67	24.2%	34.0%	35.3%
	248s	55.33	48.2%	30.7%	12.7%
	494s	57.00	39.8%	14.0%	21.1%
	742s	48.67	53.4%	53.4%	37.0%

mitochondrial structures to be more easily matched between consecutive frames, while the treated cells were acquired with larger time intervals between the image frames. This was done to validate the ability of MEL to detect mitochondrial events in conditions characterised by decreased spatio-temporal resolution.

Figures 6.4-6.7A show the maximum intensity projection (MIP) of the MEL output for every fourth frame in the time-lapse sequence. Column B shows the MIP of a region of interest (ROI), indicated by a white square in column A, for all the frames in the time-lapse sequence. The ROI was specifically chosen to highlight the detection and visualisation of mitochondrial events. This column shows both Frame 1 and Frame 2 that was used by MEL as well as the generated output frame. Column C shows a selection of the MEL ROI frames in three dimensions. Column D shows a graphical summary of some of the data extracted by MEL for each frame pair. This includes the five-frame simple moving average (SMA) of the number of mitochondrial events that occur, the fission/fusion ratio, the total number of mitochondrial structures, as well as the average volume of these structures. Lastly, Figure 6.8 shows the 3D visualisation of the four cells with the mitochondrial event locations overlaid on the mitochondrial network. All sample visualisation and ROI selection was accomplished using the virtual reality based biological sample processing system that was developed in Chapters 2, 3 and 4. The distribution of the mitochondrial events throughout the sample can only be fully appreciated when visualised in virtual reality since this provides an unambiguous interpretation of the location of the events within the complex mitochondrial network.

Finally, the average percentage of structures that will undergo fission, fusion and depolarisation in every fourth frame in the time-lapse sequence (shown in column A of Figures 6.4-6.7) was calculated, and the results are summarised in Table 6.1.

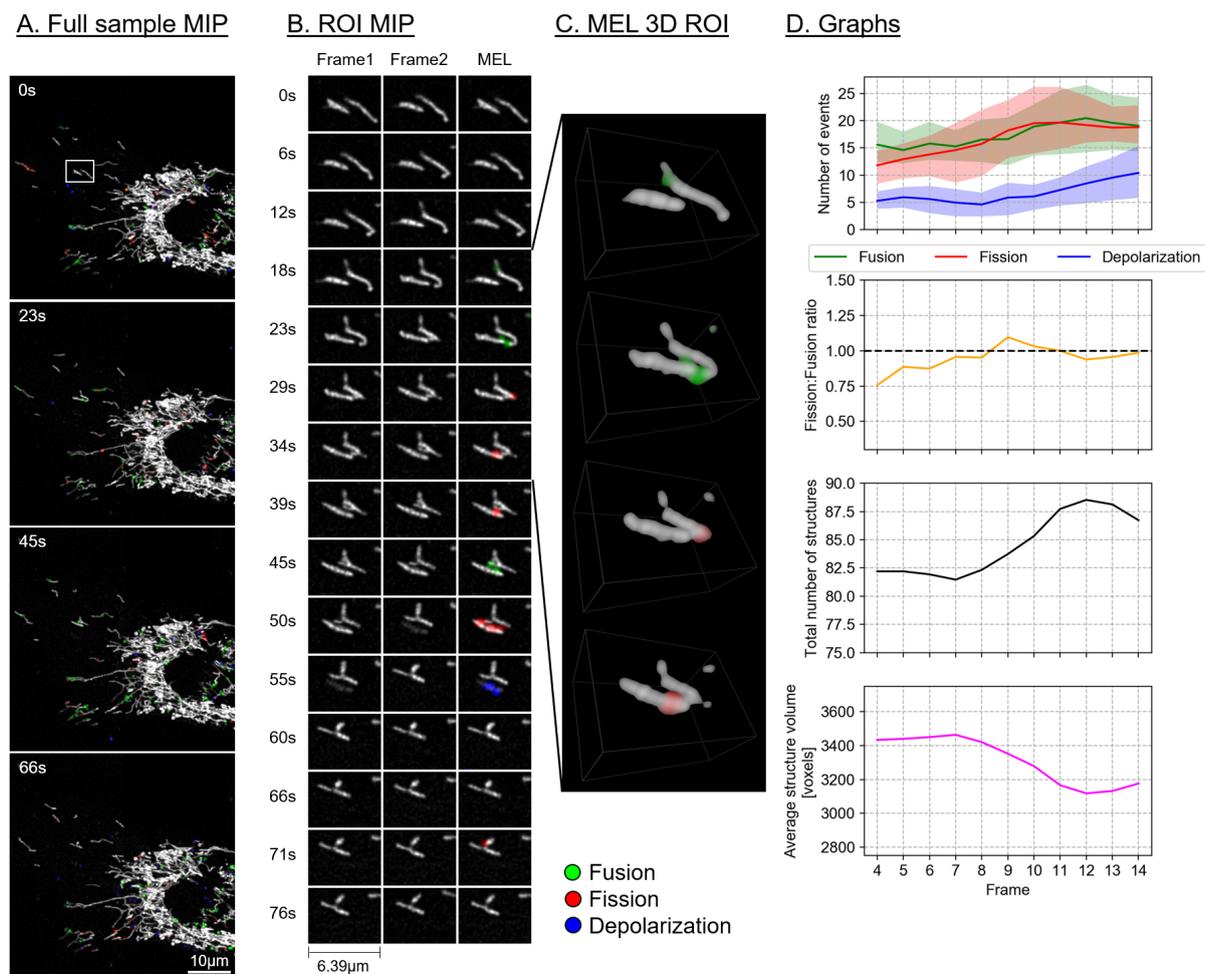


Figure 6.4: Control sample 1. A. Maximum intensity projection (MIP) of the MEL events overlaid on the entire sample image of every fourth frame in the time-lapse sequence. B. A region of interest (ROI) selection, indicated by the white square in column A. Frame 1 matches the time indicated, Frame 2 shows the subsequent time step, MEL shows the detected mitochondrial events overlaid on Frame 1. C. A selection of the ROI frames in column B, visualised in 3D using volume rendering. D. The five-frame simple moving average of the values calculated by MEL. For the number of events, the area plots indicate the maximum and the minimum that MEL detected for different pre-processing parameters.

6.4 Discussion

We have described a novel method, MEL, which localises the key mitochondrial fission, fusion and depolarisation events based on a time-lapse image sequence. We provide evidence that, when using MEL, the relationship between fission and fusion can be described quantitatively. This allows the assessment of whether a cellular system is in mitochondrial fission/fusion equilibrium and whether mitochondrial depolarisation events are kept within a physiological range. We also visualise where and when mitochondrial events, specifically fission, fusion and depolarisation, take place in the cell. Given the regional complexity of cells, which may distinctively change in disease, with a particular distribution profile of organelles, mitochondria, ATP generation and vesicle dynamics associated with mitochondrial quality control, MEL may be used to dissect localised regions of dys-

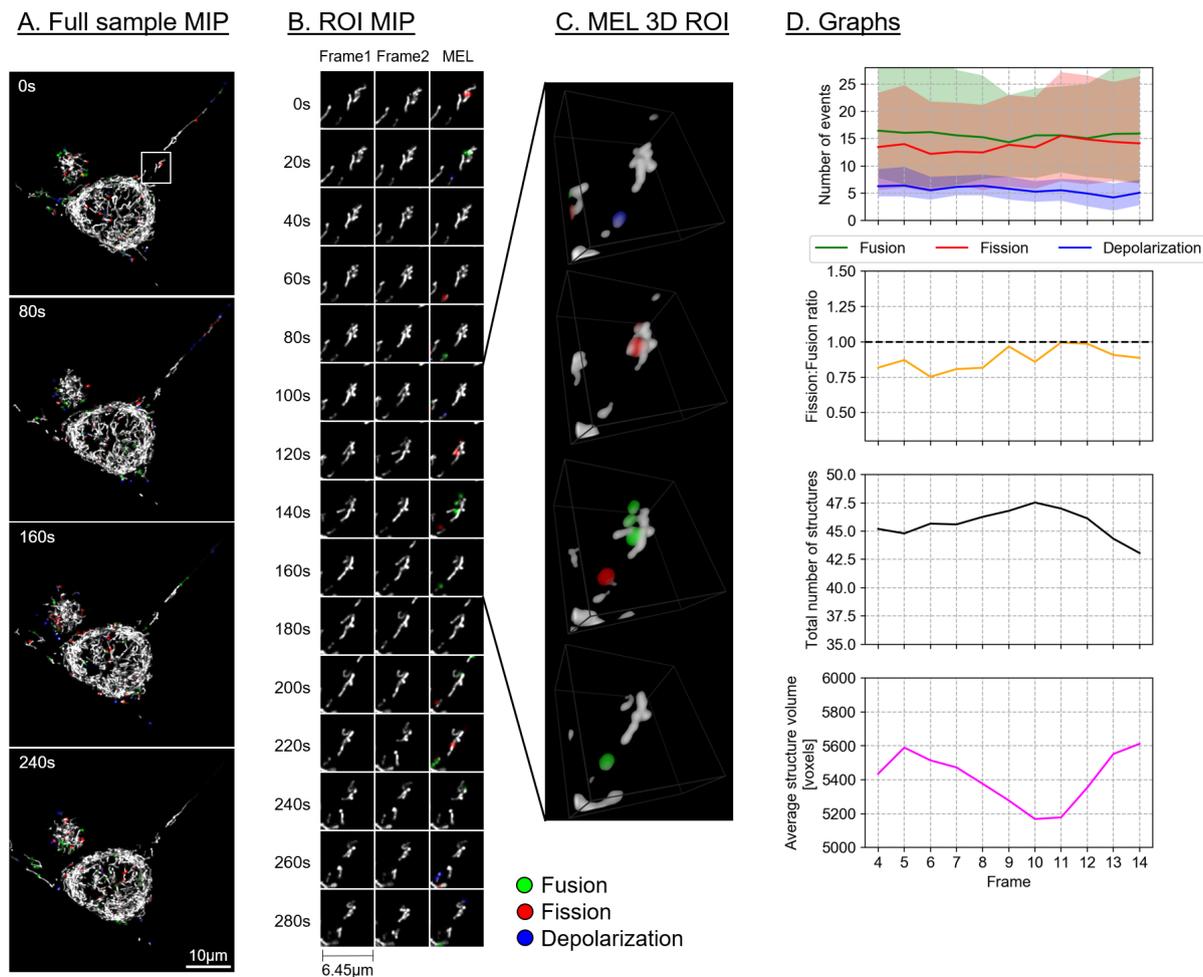


Figure 6.5: Control sample 2. Description as for Figure 6.4.

function or molecular defect precisely [128, 129, 130, 131, 132]. Moreover, when applied in scenarios where particularly mitochondrial pathology is implicated, such as in neurodegeneration or mitophagy, precise mitochondrial event localisation as provided by MEL may be of benefit [133]. Finally, due to the relatively simple data generation approach, a role for diagnostics, especially in the context of high-throughput imaging approaches, may be envisaged [134].

MEL allows the automatic identification and subsequent visualisation of mitochondrial events over the course of time-lapse sample acquisition. This allows the dynamics of mitochondrial activity to be studied. Data extracted by applying MEL, as shown in column D in Figures 6.4-6.7 and Table 6.1, reveal that the control cells are characterised by the maintenance of a fission/fusion equilibrium, with limited occurrence of depolarisation events throughout the time-lapse sequence. In contrast, HCQ treated cells exhibit much greater fluctuations in the fission/fusion equilibrium as well as a greater extent of depolarisation.

From a morphological point of view, we observe a widely spread, highly interconnected mitochondrial network for the control samples (column A in Figures 6.4 and 6.5). We note that the peripheral cell regions are characterised by a loosely defined network morphology, with few denser structures that are localised more centrally. We further notice that both fission and fusion events localise in both central and peripheral regions, in what appears to be a random fashion, with no particular focus regions or hotspots of occurrence. In

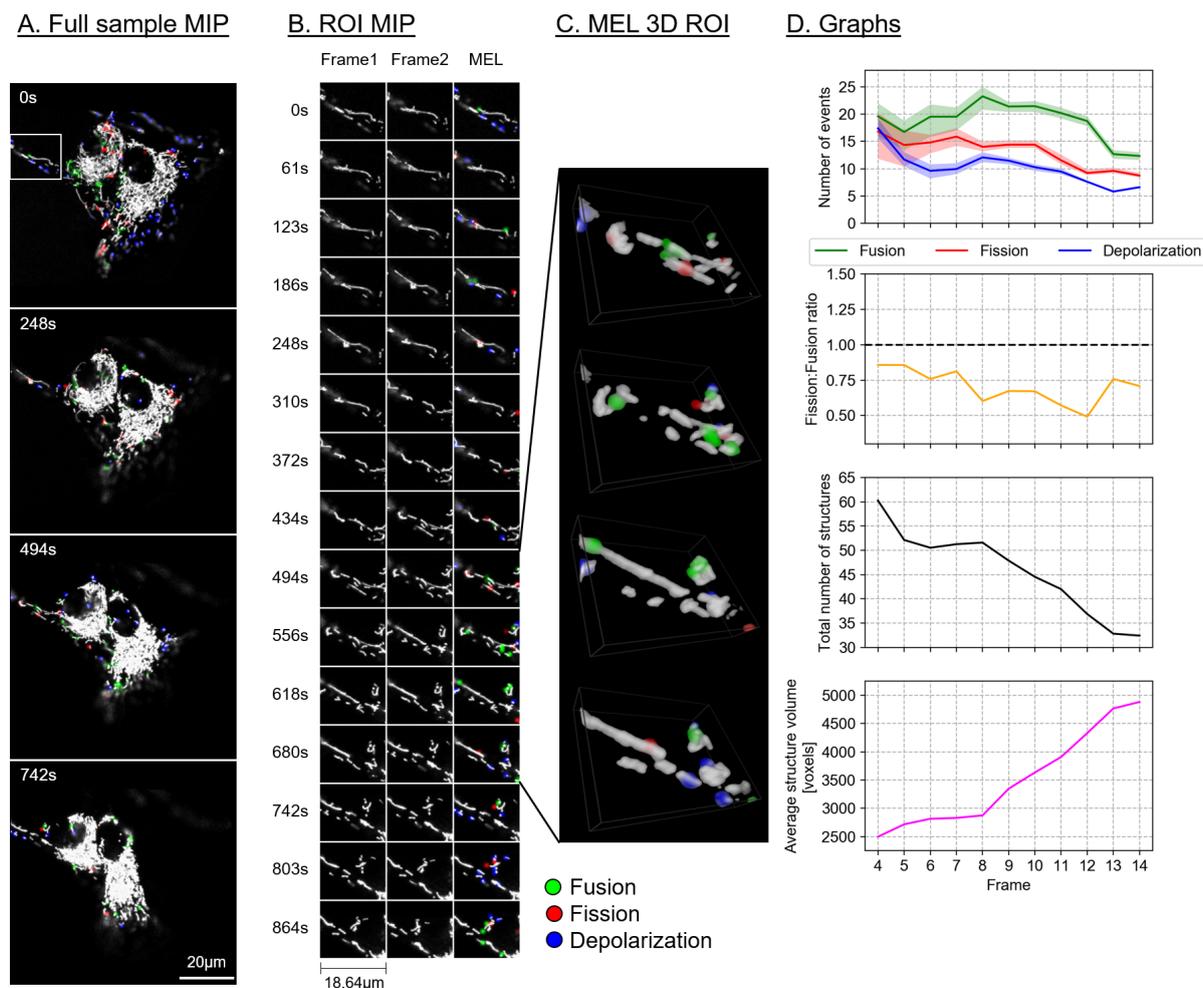


Figure 6.6: HCQ treated sample 1. Description as for Figure 6.4.

contrast, the HCQ-treated samples display a very different and distinct event profile for all mitochondrial events. Most notably, a strong occurrence of depolarisation (blue) was localised primarily in the peripheral regions, a pattern that is maintained throughout the time of acquisition (Figures 6.6 and 6.7). A bias toward fusion events (green) was also seen as the network morphology becomes more densely networked, indicating mitochondrial injury. This results in fewer mitochondrial structures, each with a higher average volume. This is, to our knowledge, the first time that depolarisation events could be localised in this manner and precision, within the context of fission and fusion events.

In general, there is very little change in the overall mitochondrial network pattern of the control samples (column A in Figures 6.4 and 6.5), even when taking the shorter interval between samples into account. In comparison, the mitochondrial network of the HCQ-treated samples (column A in Figures 6.6 and 6.7) initially appears more loosely networked and, as time progresses, is seen to become denser and centred with swollen mitochondrial morphology. This accounts for the decrease in the mitochondrial events, as well as the more consistent results for different pre-processing parameters in later frames of the time-lapse sequence.

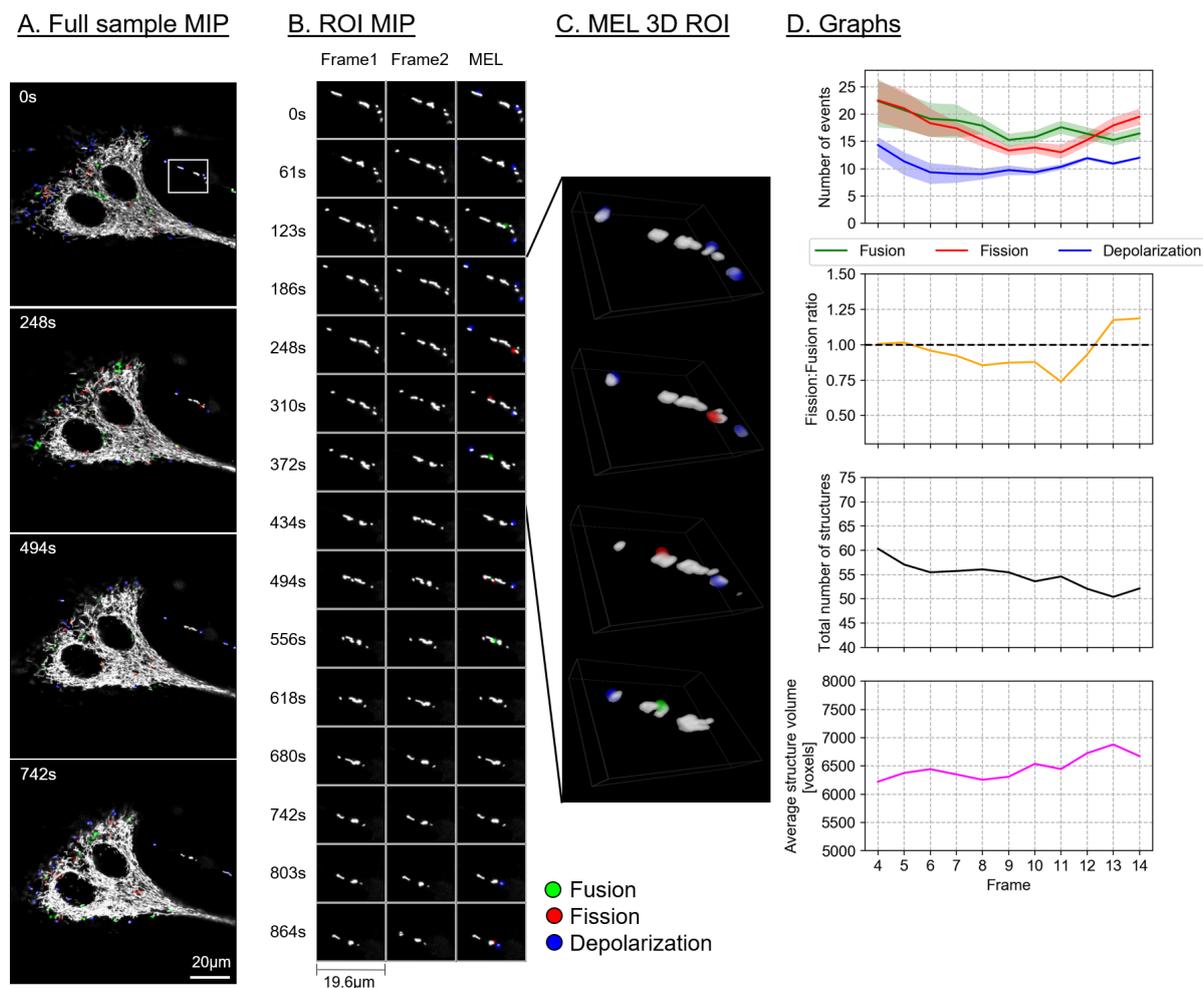


Figure 6.7: HCQ treated sample 2. Description as for Figure 6.4.

6.4.0.1 Fission and fusion events

Most the fission and fusion events were observed around the edges of the mitochondrial network, with very few events detected in the strongly networked areas. Previous work has suggested that the density of the mitochondrial network influences fission and fusion events [117].

Using Figure 6.4B and C as an example, it was often observed that the same structures alternate between fission and fusion events (between 23-50s), with possible eventual depolarisation of one of the structures (55s). This scenario is supported by previous work [119, 135], and is sometimes referred to as ‘kiss-and-run’, albeit without the 3D quantitative approach and the depolarisation context.

6.4.0.2 Depolarisation events

To our knowledge, MEL is the first approach that allows the automatic localisation of depolarisation events. This is specifically significant since depolarisation of the entire mitochondrial network has been associated with caspase-3 activation, the execution of apoptotic cell death [136] and generally demarcates the point of no return (PONR) for apoptosis [137]. Therefore, MEL may be of particular value in the quantitative assessment of cell death onset, associated with a wide range of diseases.

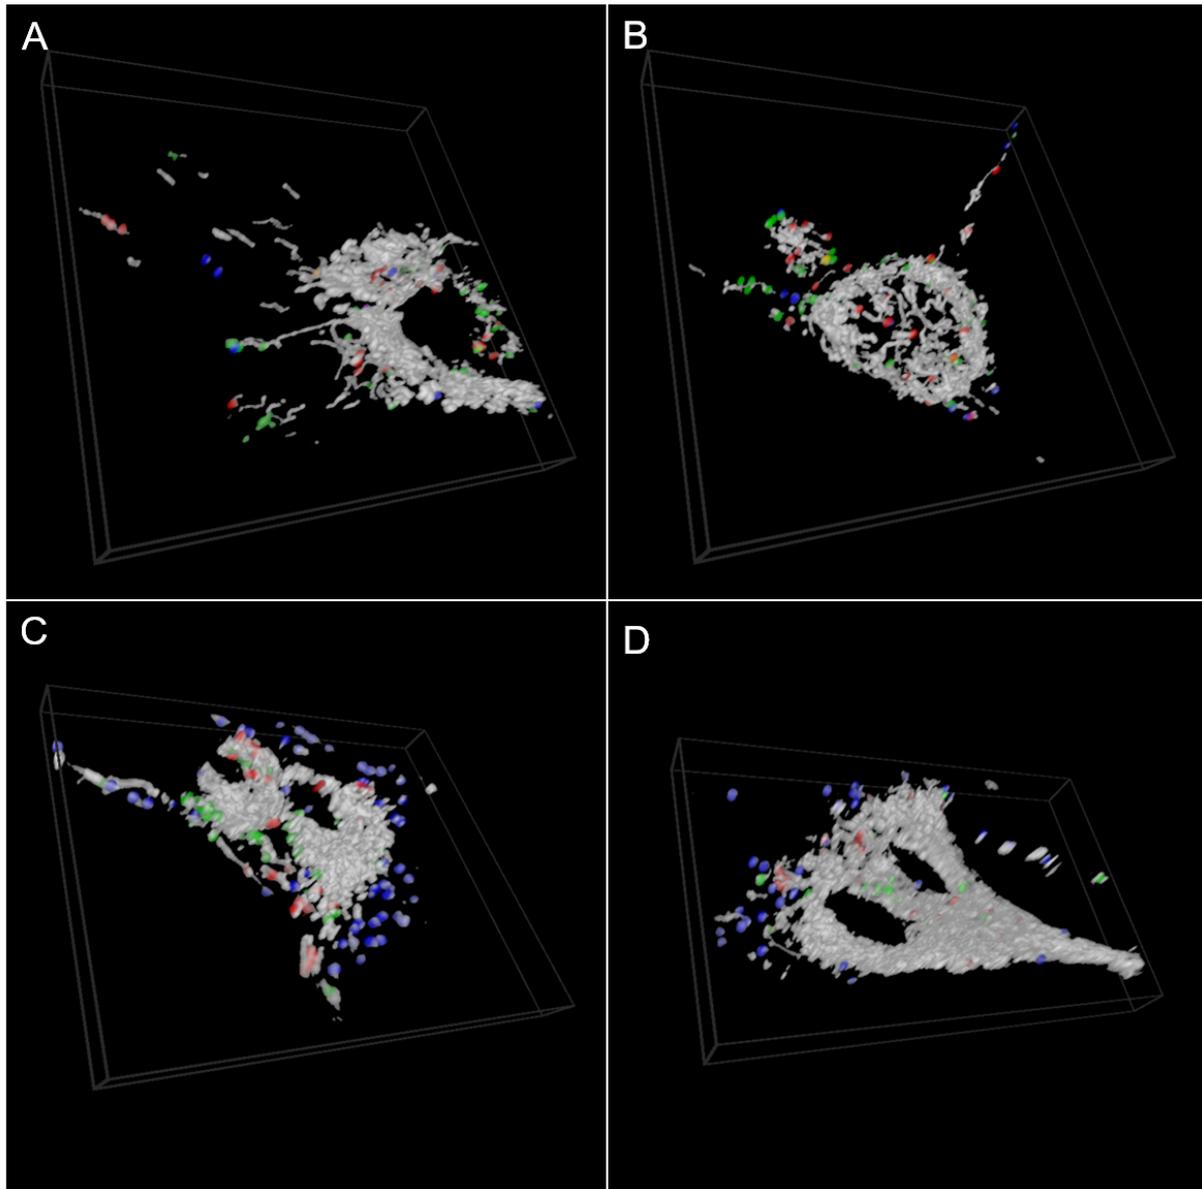


Figure 6.8: 3D visualisation of MEL event showing the first frame in the time-lapse sequence of the samples in Figures 6.4-6.7. A. Control sample 1. B. Control sample 2. C. HCQ treated sample 1. D. HCQ treated sample 2.

Depolarisation was observed to be most abundant in the HCQ treated samples, with the majority of the depolarisation events localised in mitochondria that become isolated from the main mitochondrial network and have a small average volume. This was usually observed to occur in the cell periphery (column A in Figures 6.6 and 6.7).

6.4.1 Limitations of MEL

The samples that were investigated here represent a variety of intervals between image frames as well as resolutions. From the analysis it seems that MEL is robust to these variations when the movement of mitochondrial structures between two frames is small enough.

One limitation of using binarised images is that two mitochondria passing each other, due to the motility of mitochondria, will be labelled as a fusion event at one point in the

Table 6.2: The tuneable parameters used by MEL, and the values used for experimentation.

Tuneable parameter	Value
Noise volume	20 voxels
Gaussian blur (x-y)	$\sigma = 1.5$
Gaussian blur (z)	$\sigma = 0.5$
Volume overlap threshold	$\sigma = 1\%$
Relative percentage threshold	50%
Distance threshold	20px

time-lapse and as a fission event at a later point [117]. This is a limitation of the resolving power of the microscope used, however. Since our analysis is based on 3D z-stack data, the false detection of these events is reduced.

MEL does have tuneable parameters that affect how it performs. These are the pre-processing parameters as well the percentage and distance thresholds, which are used to remove false matches. The tuneable parameters are summarised in Table 6.2. For the investigation described in our Methods, these parameters were selected based on empirical observation. While we believe that the chosen values for these parameters generalise well, an automated statistical method could also be employed to determine them based on the z-stack under analysis. Such a strategy was used, for example, for an automated image analysis algorithm that determines the optimal filtering parameters to produce the best 3D quantification of mitochondrial morphology [138].

One challenge with MEL is that it is dependent on the quality of the binarised frames (B_1 and B_2). The ability of Otsu thresholding to perform well is in turn dependent on the quality of the input images. It is for this reason that we first normalise the images before applying Otsu thresholding to ensure it produces the most consistent results. However, if the input image quality is poor, a small adjustment in pre-processing parameters can have a large impact on the resulting binarised image. The consequence of this is that some structures could be joined together for some pre-processing parameter choices and be binarised as separate structures with others. The result of this can be seen for our control samples, where there is a large difference between the minimum and the maximum number of mitochondrial events for different pre-processing parameters (shown in Figures 6.4D and 6.5D). This can be ascribed to the very limited dynamic range in the input image intensity, which is a consequence of the acquisition parameters that were chosen to ensure that longer time-lapses in 3D can be acquired while limiting the amount of bleaching and phototoxicity. Hence, well controlled acquisition parameters and an awareness of the dynamic range will benefit the application of MEL.

6.5 Conclusion

Mitochondrial events, particularly fission, fusion and depolarisation, play a central role in cellular homeostasis, function and viability. However, the precise localisation of these events within a cell as well as the quantitative characterisation thereof remains challenging.

We have presented MEL (mitochondrial event localiser), an approach we have developed to automatically determine the three-dimensional location of fission, fusion and depolarisation events in a sample from a three-dimensional time-lapse sequence. The

number of mitochondrial structures, their combined and average volume, as well as the number of events at each frame in the time-lapse sequence can subsequently be calculated to provide a quantitative assessment.

To demonstrate its performance, we have applied MEL to three-dimensional time-lapse sequences of two control samples and two hydroxychloroquine sulphate treated samples with varying time intervals and image resolutions. The analysis confirmed that the fission/fusion equilibrium is maintained throughout the time-lapse of the control cells, as is commonly expected for healthy cells. This is not the case for the treated samples, both of which showed an initial bias toward fusion, causing the mitochondrial network to become denser and more perinuclear localised.

It is hoped that the automatic localisation and quantification of mitochondrial events offered by MEL can support research into mitochondrial function in health and disease, which may consequently permit new diagnosis and treatment strategies to be implemented for diseases associated with mitochondrial dysfunction.

Chapter 7

Conclusion and future work

Modern fluorescence microscopy is a powerful tool for molecular biologists, particularly neurobiologists, which provides insights into the inner working of sub-cellular structures. Existing methods of visualisation and analysis can, in some cases, lead to ambiguity or false positives when performing a statistical analysis. To overcome these limitations, we have presented a multifaceted virtual reality based system for the three-dimensional visualisation and analysis of fluorescence microscopy samples.

The virtual reality system was applied to the calculation and analysis of colocalisation of fluorescence channels in a three-dimensional biological sample. It was demonstrated how the sample can be manipulated and regions of interest can be selected by using either a fully-immersive hand tracking system or a traditional handheld controller. A set of user trials showed that for most tasks the users were most productive, to a statistically significant degree, when using a hand tracking interface. In the subjective assessment most users indicated that they perceived the hand tracking to be the most intuitive but most users also indicated that they were under the impression that they were more productive when using the handheld controller. This negative subjective bias towards the productivity of the hand tracking was ascribed to current inaccuracies in the hand tracking process, which often led to frustration. It is likely that, as hand tracking technology improves, this bias will diminish and hand tracking interfaces could become the preferred option for interaction in the virtual reality environment.

Subsequently, a thorough investigation of colocalisation analysis was performed by comparing the proposed VR system to the two-dimensional maximum intensity projection analysis currently in use. Using neuronal cells, the dissociation between Tau and the microtubulin network in an *in vitro* model of autophagy dysfunction was considered, since this mimics a key pathology in Alzheimer's disease. The change in colocalisation between the two structures at a control, after 6 hour and again after 24 hour of neuronal injury, in different regions of the cell were considered. Statistically significant decreases in colocalisation over time were observed when using the three-dimensional analysis, indicating that the dissociation of Tau and microtubulin may be an early event in neuronal injury prior to the onset of cell death. This tendency could not be consistently observed when using the two-dimensional analysis. In fact, the analysis based on the two-dimensional Pearson's correlation coefficient in some cases indicated increasing colocalisation with time. Furthermore, by using the region of interest selection methods that was implemented as part of the virtual reality system, it could be shown how very small structures, such as protein aggregates and microtubulin strands, can be isolated and analysed independently and with high precision.

Recognising the need for a robust visual colocalisation analysis method that biologists

can use to gain a better understanding of the association between molecular structures within a cell, we developed a new quantitative visual analysis method that offers a spatial visualisation of colocalisation in fluorescence-based micrographs using a colourmap. Regression adjusted colocalisation colour mapping (RACC) was designed to meet three objectives: firstly to highlight regions within a cell that have a strong positive correlation between two fluorescence channels, secondly to highlight colocalised regions that have greater combined fluorescence intensities and thirdly to suppress coincidentally colocalised voxels from the visualisation. These objectives were achieved by applying Deming regression in conjunction with a geometric projection of the fluorescence intensities while using automatically calculated thresholds to remove outliers. We demonstrated and validated RACC by visualising both synthetic and biological data. These visualisations demonstrated how RACC can offer an improved understanding, compared to existing methods, of the nature of the colocalisation within a sample by visualising only truly colocalised regions while offering a qualitative measure of the correlation between fluorescence channels in different regions of a sample.

Since it is becoming increasingly clear that mitochondrial events, particularly fission, fusion and depolarisation, play a central role in cellular homeostasis, function and viability, the final part of the dissertation proposed an approach allowing the automatic determination of the three-dimensional location of mitochondrial events within a fluorescence based time-lapse sequence. The number of mitochondrial structures, their combined and average volume, as well as the number of events at each frame in the time-lapse sequence can be calculated to allow a quantitative assessment. This method, referred to as the mitochondrial event localiser (MEL), was applied to time-lapse sequences of two control samples and two HCQ treated samples with varying time intervals and image resolutions. The analysis confirmed that the fission/fusion equilibrium is maintained throughout the time-lapse sequence, as is commonly expected for healthy cells. This is not the case for the treated samples, both of which showed an initial bias toward fusion, causing the mitochondrial network to become denser and more perinuclear localised.

We conclude that virtual reality offers an attractive and powerful means of visualisation for microscopy data and, by using 3D analysis, samples can be interrogated and assessed with greater precision, thereby better exploiting the potential of fluorescence-based image analysis in biomedical research.

Future work

The work presented in this dissertation may serve as a basis upon which further analysis tools may be developed, associated with enhanced intracellular, 3D based precision analysis. This may include the following:

- The adaption of the VR system that incorporate new input interfaces which combine hand tracking with handheld controllers, such as the Oculus Touch or HTC Vive controllers. This could enhance both the efficiency and the control of the analysis, as well as more intuitive interactions with the interface. Such an approach may further enhance sample control and navigation, thereby improving subsequent analysis.
- Using the developed VR system as a point of departure, specifically designed teaching tools may be developed, focusing on cellular function and sub-cellular component interactions in an interactive way.

- The validation of MEL using a variety of treated cells and extending MEL to geometrically indicate in which region of the cell the different mitochondrial events localise.
- The training of neural networks to estimate the mitochondrial event locations from a single z-stack, without the need of a time-lapse sequence. This can be achieved by using several MEL generated input-output pairs as the ground truth data for neural network training. The most promising network architecture to achieve this is the conditional generative adversarial networks [139, 140], such as the pix2pix network [141]. If successful, this would greatly enhance the ability to perform high-throughput image analysis for drug screening. Indeed, current trends indicate the strength of utilising high-throughput image analysis tools to perform functional cell analysis to screen for novel drugs and to exclude cell toxicity [142].
- The implementation of MEL using key model systems of clinical relevance will allow the building of standards and databases for fission/fusion equilibrium under both control and diseased condition for the candidate model systems. This will allow the screening of drugs in the context of neurodegenerative diseases and cancer, where either mitochondrial dynamics are aimed to be maintained and preserved (for neurodegenerative diseases), or effectively disrupted (cancer). This approach may benefit from being coupled to the detection of other critical parameters associated with cellular function, such as caspases activity [143].
- The adaptation of the MEL process to calculate several other commonly utilised quantitative mitochondrial morphological features in three dimensions may be an additional future avenue of interest. These features, which include the form factor, aspect ratio and the perimeter, are typically calculated only in two dimensions. A system that consolidates such features within a single analysis tool, and also extends the analysis to the equivalent three-dimensional features, will contribute to the generation of new knowledge and insights into mitochondrial function. Since previous work has shown that these morphological features can predict fission and fusion events [117, 122], it may also be possible to train a neural network that receives these features as inputs and produces a prediction of the cell's fate.

Appendices

Appendix A

Direct volume rendering implementation: Texture-based approach

One of the main methods to implement direct volume rendering (DVR) is known as texture-based volume rendering, also called texture slicing. We introduced textured-based volume rendering in Section 2.3.2.3, in this appendix we will expand on this and consider more detailed implementation details. We will only discuss the viewer-aligned slices variant of texture slicing since it produces a clearer visualisation which is better suited for virtual reality.

Our implementation has been adapted from the algorithms described by Movania [22] and Engel *et al.* [20].

A.1 Overview

Texture-based volume rendering is an *object-order approach*, which means that the software iterates over the object slices to create the visualisation. These slices can also be defined arbitrarily, in the form of *proxy geometry*, which can be rendered instead of the fluorescence microscopy z-stack images which stores the volume data. In the case of the viewer-aligned slices approach, this proxy geometry is defined to be perpendicular to the viewing direction. This prevents the user from being able to look in-between slices, thereby creating a more compelling visualisation. Finally, the proxy geometry is rendered using a GPU during which *trilinear interpolation* is applied over the surface of the proxy geometry slices by resampling of the volume data. These semi-transparent slices are then blended in order to create a three-dimensional (3D) visualisation of the fluorescence microscopy sample.

This process relies on the fact that the volume data can be stored as a 3D texture in GPU memory. This 3D texture is not a volumetric primitive but rather a section in memory that can be used as a look-up-table (LUT) within the fragment shader (see Section 2.2.4). The proxy geometry's vertices are therefore used as reference values in this LUT. This is illustrated in Figure A.1.

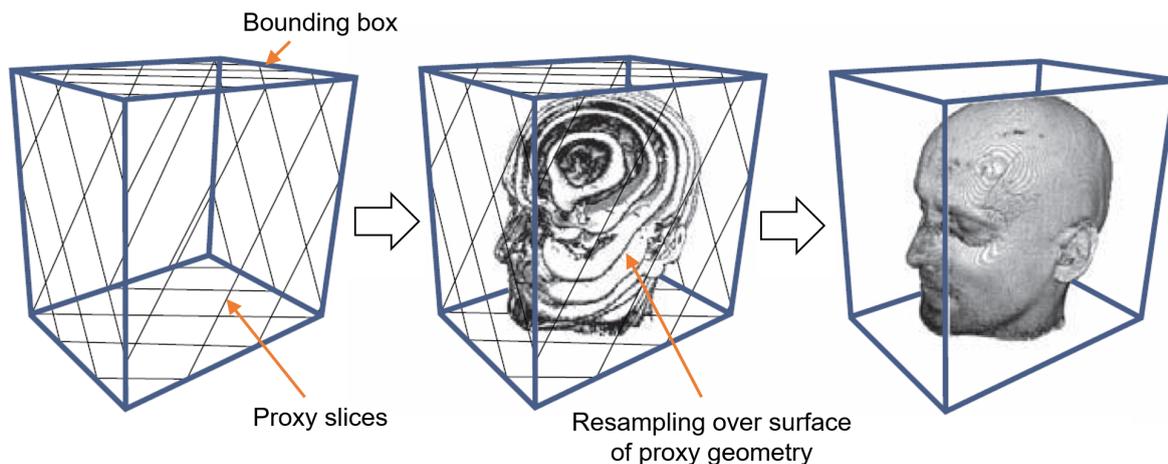


Figure A.1: A bounding box is defined around the volume data. Proxy geometry are then generated inside the box. The volume data, stored as a 3D texture, is then resampled over the surface of the proxy geometry slices to create the volume rendering. Adapted from Engel *et al.* [20]

A.2 Proxy geometry set-up

The calculation of the proxy geometry is a multi-step process. Before the proxy geometry itself can be calculated, it is necessary to define a bounding box around the volume data. The proxy geometry can then be defined as a plane that cuts through this bounding box. The resulting polygon's vertices are the locations of where the plane intersects with the edges of the bounding box. The 3D texture data is then resampled and interpolated over the surface of this polygon.

One disadvantage of the viewer-aligned slices approach is that the proxy geometry has to be recalculated whenever the viewing direction changes [20]. This could happen several times a second during an interactive visualisation. It is, therefore, imperative to use an efficient algorithm to calculate the plane-box-intersections producing the proxy geometry.

A.2.1 Defining the bounding box

The bounding box represents the limits within which the three-dimensionally reconstructed volume will be rendered and is defined as a rectangular prism that has a width, height and depth. The vertices (corners) that describe the bounding box correspond to the corners of the 3D texture which stores the volume data.

Since fluorescence microscopy data sets do not all have the same resolution and number of slices, it becomes necessary to vary the dimensions of bounding box to maintain the relative scale as to match that of the volume data. This influences some of the standard calculations, as will be shown in the following sections. Furthermore, the bounding box is defined to have its centre at the origin. Therefore, if for example the bounding box is a unit cube, for each of the x , y and z coordinates the box will be defined within the $[-0.5, 0.5]$ range.

A.2.2 Mapping fragment coordinates to texture coordinates

Texture coordinates are defined as values between 0.0 and 1.0 in all three dimensions. Therefore, in order to sample the volume data stored in a 3D texture over the surface of

the proxy geometry, the vertices defining the proxy geometry must be mapped to valid texture coordinates. This mapping is performed in the vertex shader, to ensure that the fragment shader receives the correctly interpolated texture coordinates (see Section 2.2.4).

Since we defined the bounding box, which contains the proxy geometry, to be centred around the origin, we can map the texture coordinates to the range $[0.0, 1.0]$ using the following equation:

$$\text{Texture coordinate} = \left(\frac{x_i + \frac{W}{2}}{W}, 1 - \frac{y_i + \frac{H}{2}}{H}, \frac{z_i + \frac{D}{2}}{D} \right) \quad (\text{A.1})$$

In Equation A.1, the dimensions of the bounding box is defined by W , H and D for the width, height, and depth, respectively. The position of the i^{th} vertex on the proxy geometry is defined by the coordinate point (x_i, y_i, z_i) . The resulting texture coordinate is what is passed to the fragment shader. The y -coordinate is inverted in order to maintain the correct orientation of the sample, since y -coordinates increase from the top to the bottom in images, but from the bottom to the top in textures.

A.2.3 Calculating the proxy geometry

For the viewer-aligned slices approach, calculating the intersection points of the slice planes with the bounding box, producing the proxy geometry, is a more complicated task than in the case of object-aligned slicing, where all the slices are rectangular and aligned with the axes of the box. This is because the viewer-aligned slices are arbitrarily orientated planes that intersect the bounding box depending on the position of the observer. This results in a proxy geometry polygon consisting of between three and six vertices as are illustrated in Figure A.2.

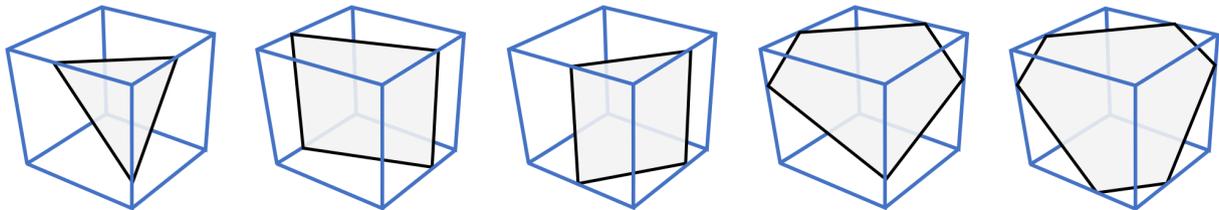


Figure A.2: An illustration of several arbitrary planes slicing through the bounding box producing proxy geometry polygons consisting of between three and six vertices. Adapted from Engel *et al.* [20].

It is, however, not computationally efficient to change the number of vertices of each proxy geometry polygon every time they are updated after a change in viewing direction. We will, therefore, always represent them using a six vertex polygon. When the proxy geometry polygon has fewer than six vertices, the slicing algorithm will duplicate one or more vertices. Two identical vertices imply that the triangular primitive will have an edge of zero length. This results in a degenerate triangle which does not produce any fragments during the rasterisation stage [20].

The rendering of quads and polygons have been deprecated in modern OpenGL (version 3.3+), and the only primitives still supported are points, lines and triangles (including variations of these, such as a line or a triangle strip). To form the slicing polygon a variation of the triangle geometric primitive, known as a *triangle fan*, will be used. An example of a six vertex triangle fan is shown in Figure A.3. The first vertex is always starting

vertex. Thereafter, every two adjacent vertices form a triangle with the first vertex. By using a triangle fan instead of separate triangles, the memory requirements are halved. From Figure A.3 is clear that only six vertices need to be stored in memory, however, to render four separate triangles, a total of 12 vertices (3 vertices times 4 triangles) must be stored.

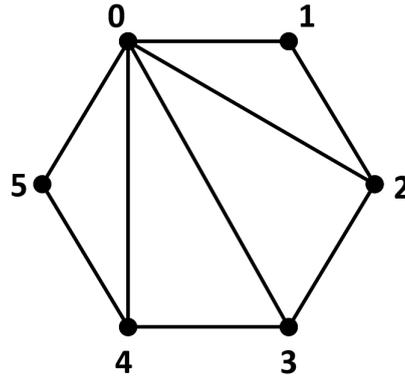


Figure A.3: An illustration of a six vertex triangle fan. Vertex 0 is the starting vertex, to which all other vertices connect.

Calculating the plane-box-intersections

The following calculations of the plane-box-intersections are based on the derivation given by Engel *et al.* [20].

It is easy to determine the intersection of a slice plane with the edge of the bounding box if the plane is given in the Hessian normal form,

$$\mathbf{n}_p \cdot \mathbf{x} = d \quad (\text{A.2})$$

where \mathbf{n}_p denotes the normal vector of the plane, \mathbf{x} denotes a location vector of any point on the plane, and d is the distance between the plane and the origin. In the case of viewer-aligned slices, the normal vector \mathbf{n}_p is also the camera's viewing direction.

An edge ($E_{i \rightarrow j}$) between two vertices (\mathbf{v}_i and \mathbf{v}_j) of the bounding box can be described by:

$$E_{i \rightarrow j} : \quad \mathbf{x} = \mathbf{v}_i + \lambda(\mathbf{v}_j - \mathbf{v}_i) \quad (\text{A.3})$$

$$= \mathbf{v}_i + \lambda \mathbf{e}_{i \rightarrow j} \quad \text{with } \lambda \in [0, 1]. \quad (\text{A.4})$$

In Equation A.4, $\mathbf{e}_{i \rightarrow j}$ is the vector along the edge from vertex i to j and will have unit length in the case of a unit cube. This is, however, not generally the case. Furthermore, the position, as a fraction of the edge length, where the plane intersects the edge in question is denoted by the *intersection parameter* λ . This intersection fraction is calculated by rearranging Equation A.4 to become

$$\lambda = \frac{d - \langle \mathbf{n}_p \cdot \mathbf{v}_i \rangle}{\langle \mathbf{n}_p \cdot \mathbf{e}_{i \rightarrow j} \rangle} \quad (\text{A.5})$$

In Equation A.5 the denominator becomes zero, and the fraction undefined, when the edge is co-planar with the slice plane. In this case the intersection is simply ignored. A

valid intersection occurs only when λ is in the range $[0, 1]$, otherwise the plane does not intersect the edge under consideration.

The most difficult part of performing the intersection calculations is to maintain a valid ordering of the intersection vertices forming each proxy geometry polygon. In particular, the resulting vertices must form a valid triangle fan. The algorithm used to achieve this will be described by using an example based on Figure A.4. In this figure the vertex that is the closest to the camera is vertex \mathbf{v}_0 and vertex \mathbf{v}_7 lies on the opposite corner of the bounding box's diagonal, making it the furthest vertex from the camera.

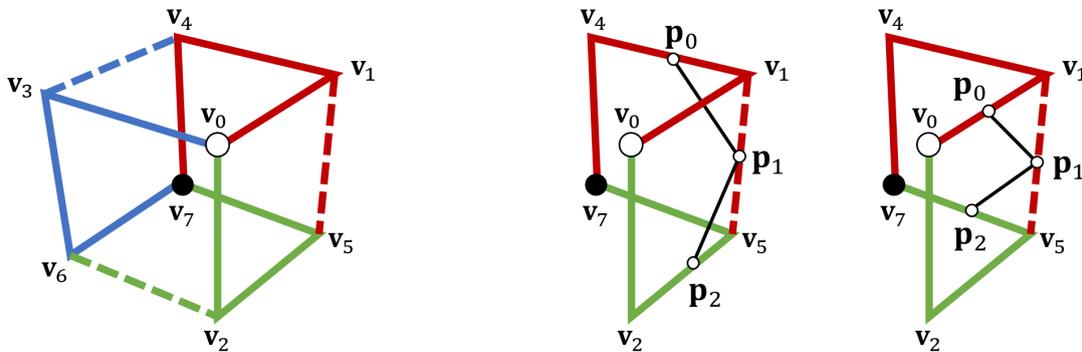


Figure A.4: Left: The vertex numbering that is used for the bounding box. There are always three independent paths from the front vertex \mathbf{v}_0 to the back vertex \mathbf{v}_7 , as marked by the solid lines. Right: The plane intersection with the dotted line will always occur between the intersection with two solid lines and must be inserted accordingly. Adapted from Engel *et al.* [20].

With \mathbf{v}_0 the front (closest) vertex and \mathbf{v}_7 the back (furthest) vertex, there are exactly three independent paths to move from \mathbf{v}_0 to \mathbf{v}_7 via the edges, as is indicated by the solid lines in Figure A.4. By independent we mean that the paths do not share any vertices except for the start and end vertices. Each of these paths consists of three edges $\{E_1, E_2, E_3\}$. For example, in the case of the red path, $E_1 = E_{0 \rightarrow 1}$, $E_2 = E_{1 \rightarrow 4}$ and $E_3 = E_{4 \rightarrow 7}$.

Now consider a viewer-aligned slice plane that sweeps through the box from the front to the back. The first vertex that is touched is \mathbf{v}_0 . Before this, there is no valid intersection with the box. The last vertex the plane touches before leaving the box is \mathbf{v}_7 . After this, there is also no valid intersection with the box. In between these extremes, the plane will always have exactly one unique intersection with each of the three solid paths. Therefore, in the case that the plane polygon has only three vertices, they will be three points on these three paths. For this reason we can calculate three of the possible six intersection points \mathbf{p}_i by checking for intersections with each of the edges on each path:

$$\begin{aligned} \mathbf{p}_0 &= \text{Intersection with } E_{0 \rightarrow 1} \text{ or } E_{1 \rightarrow 4} \text{ or } E_{4 \rightarrow 7}, \\ \mathbf{p}_2 &= \text{Intersection with } E_{0 \rightarrow 2} \text{ or } E_{2 \rightarrow 5} \text{ or } E_{5 \rightarrow 7}, \\ \mathbf{p}_4 &= \text{Intersection with } E_{0 \rightarrow 3} \text{ or } E_{3 \rightarrow 6} \text{ or } E_{6 \rightarrow 7}. \end{aligned}$$

Now the remaining three possible intersections need to be considered. These will lie on the dotted edges shown in Figure A.4. If these intersections are valid, as determined from the respective λ values, they should be inserted in-between the points already found on the paths in order to form a valid polygon. Taking the case of the red path as an example, if an intersection on the dotted edge $E_{1 \rightarrow 5}$ exist, the point must be inserted

between the intersections with the red and green path, as shown in Figure A.4. If there is no intersection, the point is simply set to \mathbf{p}_0 (as was discussed in the beginning of Section A.2.3). Therefore, the remaining three points are set as follows:

$$\begin{aligned}\mathbf{p}_1 &= \text{Intersection with } E_{1 \rightarrow 5}, \text{ otherwise } \mathbf{p}_0, \\ \mathbf{p}_3 &= \text{Intersection with } E_{2 \rightarrow 6}, \text{ otherwise } \mathbf{p}_0, \\ \mathbf{p}_5 &= \text{Intersection with } E_{3 \rightarrow 4}, \text{ otherwise } \mathbf{p}_0.\end{aligned}$$

Now all six possible intersection points of the viewer-aligned slice plane with the bounding box are determined in the correct sequence to form a valid proxy geometry polygon. This polygon is then rendered using a triangle fan (see Figure A.3).

Lastly, these planes are generated, with their normal equal to the viewing direction vector, at fixed distances from each other throughout the bounding box to produce the volume rendering.

A.3 Conclusion

In this appendix we discussed the implementation of the texture-based direct volume rendering method. We specifically implemented the approach of using proxy geometry slices that are aligned with the viewing direction. We focused on how the vertices of the proxy geometry polygons are calculated. In Appendix B, we consider an alternative direct volume rendering method known as volume ray casting.

Appendix B

Direct volume rendering implementation: Volume ray casting

Volume ray casting is an algorithm to perform direct volume rendering (DVR). We introduced this algorithm in Section 2.3.2.2, in this appendix we will expand on the explanation and provide more detailed implementation details.

Our implementation has been adapted from the algorithms described by Movania [22] and Engel *et al.* [20].

B.1 Overview

Volume ray casting is known as an *image order* volume rendering technique, which means that the algorithm iterates over the pixels of the final rendered image to produce the visualisation of the data. Practically this is achieved by rendering a bounding box that represent the limits within which the reconstructed volume will be rendered. The vertices (corners) that describe the bounding box correspond to the corners of the 3D texture which stores the volume data. A single ray is then projected into the volume for every visible fragment on the bounding box's faces (see Sections 2.2.4 and 2.3.2.2). The volume data is then resampled at regular discrete positions along the ray, and blended together. All the calculations to perform the ray casting is done in a fragment shader, which runs on the GPU (see Section 2.2.4).

There are two different approaches to perform volume ray casting. We will refer to the first as *standard ray casting*, also called *single pass GPU ray casting*. This approach produces similar results to that of texture slicing (Appendix A). We will refer to the second method as *pseudo-isosurface ray casting*. This approach creates the impression of an isosurface, which is a surface that represents the voxels of constant fluorescence intensity which is called the isovalue. The resulting visualisation is similar to what is produced by indirect volume rendering methods where an isosurface is extracted before rendering.

B.2 Standard ray casting

Before ray casting can be performed the volume data must be stored as a 3D texture in GPU memory and a bounding box must be rendered causing fragments to be generated over the visible faces of the box during the rasterisation phase (see Section 2.2.4). Ray casting calculations are then performed on each of these fragments in order to render the

reconstructed volume. These calculations will be discussed in the following section with reference to Figure B.1.

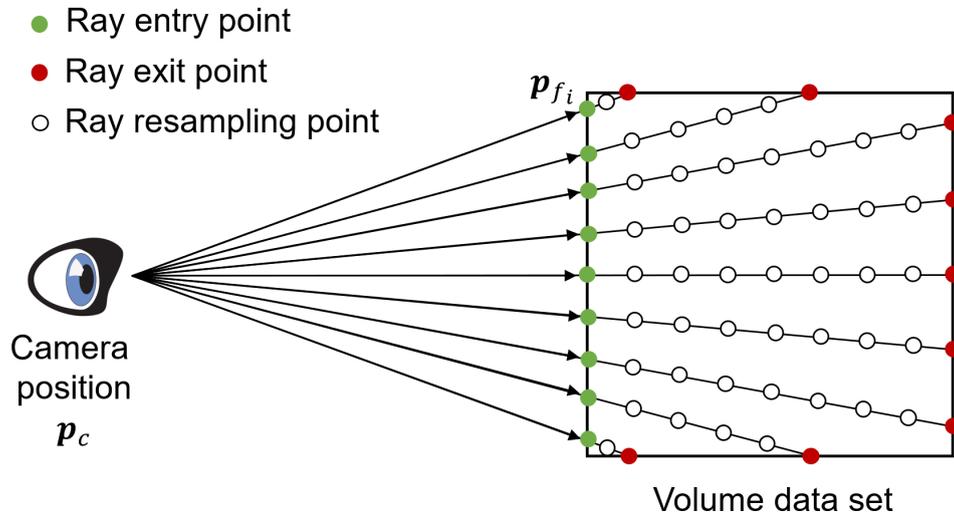


Figure B.1: Ray casting is performed by sampling the volume data set at discrete locations along the ray. Adapted from from Movania [22].

B.2.1 Rendering a bounding box

Since fluorescence microscopy data sets do not all have the same resolution and number of slices, it becomes necessary to vary the dimensions of bounding box to maintain the relative scale to match that of the volume data. This influences some of the standard calculations, as will be shown in the following sections. Furthermore, the bounding box is defined to have its centre at the origin. Therefore, if for example the bounding box is a unit cube, for each of the x , y and z coordinates the box will be defined within the $[-0.5, 0.5]$ range.

B.2.2 Mapping the fragment coordinates to texture coordinates

In the fragment shader, every fragment has a 3D vector position which can be defined as $\mathbf{p}_{f_i} = (x_i, y_i, z_i)$. The dimensions of the bounding box is also sent to the fragment shader and is defined as $\mathbf{b} = (W, H, D)$. Since the bounding box is defined to have its centre at the origin, the range of possible values for \mathbf{p} is $[-\frac{\mathbf{b}}{2}, \frac{\mathbf{b}}{2}]$. However, when a value in a 3D texture is referenced, it needs to be between 0.0 and 1.0 for all dimensions. Therefore, to map the fragment coordinates to texture coordinates \mathbf{p}_t , the following equation is used:

$$\mathbf{p}_t = \left(\frac{x_i + \frac{W}{2}}{W}, 1 - \frac{y_i + \frac{H}{2}}{H}, \frac{z_i + \frac{D}{2}}{D} \right) \quad (\text{B.1})$$

In Equation B.1, the y -coordinate is inverted in order to maintain the correct orientation of the sample, since y -coordinates increase from the top to the bottom in images, but from the bottom to the top in textures.

B.2.3 Calculating the direction of the ray

The vector that represents the direction of the ray that is cast into the volume for each fragment is called \mathbf{d}_i . \mathbf{d}_i is calculated by taking the vector difference between the camera position \mathbf{p}_c and the fragment position \mathbf{p}_f on the face of the bounding box. This ray can then be visualised as running through the bounding box, sampling the 3D texture along the ray (Figure B.1).

$$\mathbf{d}_i = \mathbf{p}_{f_i} - \mathbf{p}_c \quad (\text{B.2})$$

The vector \mathbf{d}_i in Equation B.2 is normalised after the calculation in order to create a unit direction vector that will simplify calculation that follow.

B.2.4 Calculating the step vector between samples

Now that the normalised direction of the ray has been calculated, next the vector between two contiguous samples, known as the step direction \mathbf{s} , can be calculated. A fixed number of samples N are taken between two opposite faces of the bounding box. Furthermore, since the bounding box is not necessarily a cube but an arbitrary rectangular prism, the ray's direction has to be transformed according to the dimensions of the bounding box as follows:

$$\mathbf{s}_i = \mathbf{d}_i \cdot \frac{1}{N\mathbf{b}} \quad (\text{B.3})$$

B.2.5 Sample along the ray

After determining the direction of the ray and the step vector, the volume is sampled along the ray by iterating from the fragment position on a face of the bounding box into the volume. The new 3D texture coordinate \mathbf{p}_t is recalculated for every sample step using the following recursive formula, using Equation B.1 as starting condition:

$$\mathbf{p}_t = \mathbf{p}_t + \mathbf{s}_i \quad (\text{B.4})$$

The value at the resulting texture coordinate is called C_j . If there is no data value at the sample location, a value is determined through trilinear interpolation using the surrounding samples.

B.2.5.1 Composing the final fragment colour

In Section 2.3.2.2 the recursive equations to perform front-to-back ray casting were given as:

$$C'_j = C'_{j-1} + (1 - A'_{j-1})C_j \quad (\text{B.5})$$

$$A'_j = A'_{j-1} + (1 - A'_{j-1})A_j \quad (\text{B.6})$$

with the starting conditions of $C'_0 = 0$ and $A'_0 = 0$.

These equations are used to build the final RGBA (red, green, blue, alpha) fragment colour, called C_{frag} . This is implemented by assigning the RGB part of the fragment colour to C'_j , and assigning the fragment opacity to A'_j . Since there are no opacity values in the source volume data, the value of A_i chosen as a fraction between 0 and 1 and indicates the overall opacity of the visualisation.

B.2.5.2 Stop when the ray exists the bounding box

In order to avoid calculations that do not contribute to the final image, thereby improving performance, the sampling calculations should stop as soon as the ray passes through the back of the bounding box. As mentioned before, texture coordinates must lie between 0.0 and 1.0, and since the bounding box is already mapped to the vector \mathbf{p}_t , it is possible to determine when the ray exits the bounding box when one of the dimensions of \mathbf{p}_t is no longer in the valid range. If the ray is outside, the sampling loop terminates and the next fragment is considered.

B.2.5.3 Early ray termination

Early ray termination is another way to improve the performance of the algorithm. This method stops the calculation as soon as the opacity along the partial path of the ray is greater than a predetermined threshold. This eliminates computations that do not contribute substantially to the rendered image. For our application, this threshold value was chosen as 0.9. Therefore, as soon as the alpha component of C_{frag} is greater than 0.9, the sampling loop terminates and the next fragment is considered.

B.3 Pseudo-isosurface ray casting

In order to perform pseudo-isosurface ray casting, the first few steps are the same as standard ray casting, namely those discussed in Sections B.2.1, B.2.2, B.2.3 and B.2.4. The sampling along the ray is also performed in a similar way, but produces a different visualisation as will be discussed next.

B.3.1 Sample along the ray

For every colour channel (red, green and blue), an isosurface is calculated and rendered producing a result as is shown in Figure B.2 which is an example of a fluorescence microscopy mammalian cell.

As with standard ray casting, in the pseudo-isosurface variant a ray is cast into the bounding box by iterating from the fragment position, which lies on the face of the bounding box, into the volume. The 3D texture coordinate \mathbf{p}_t is updated during every iteration of the loop using Equation B.4. The loop is terminated and the next fragment is considered as soon as the isosurface boundary has been crossed.

B.3.2 Determine isosurface boundary

In order to determine the isosurface boundary two samples are taken per colour channel, one at the previously calculated texture coordinate \mathbf{p}_t , and the other at the next sample position $\mathbf{p}_t + \mathbf{s}_i$. These two samples are then compared to the user defined isovalue. If one is greater than and the other is less than this value, a boundary has been detected. An illustration of this process is shown in Figure B.3. In order to improve the rendering output, the exact boundary within the trilinearly interpolated volume data is then determined using *bisection based refinement* which is discussed next.

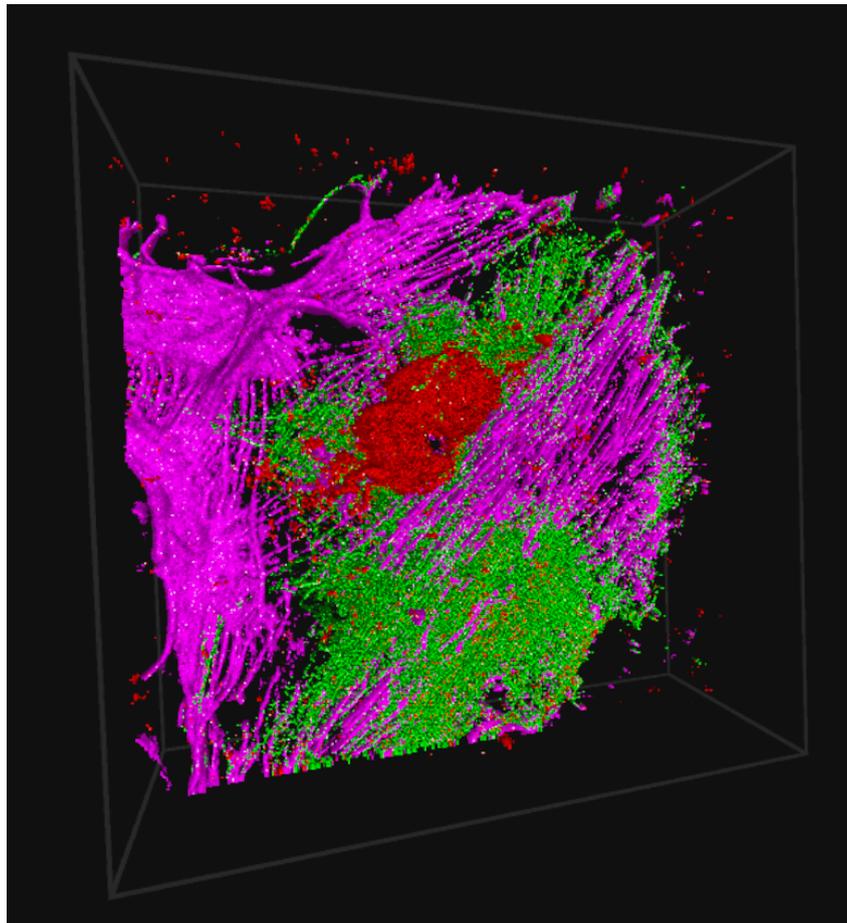


Figure B.2: An example of applying pseudo-isosurface ray casting volume rendering to a fluorescence microscopy sample.

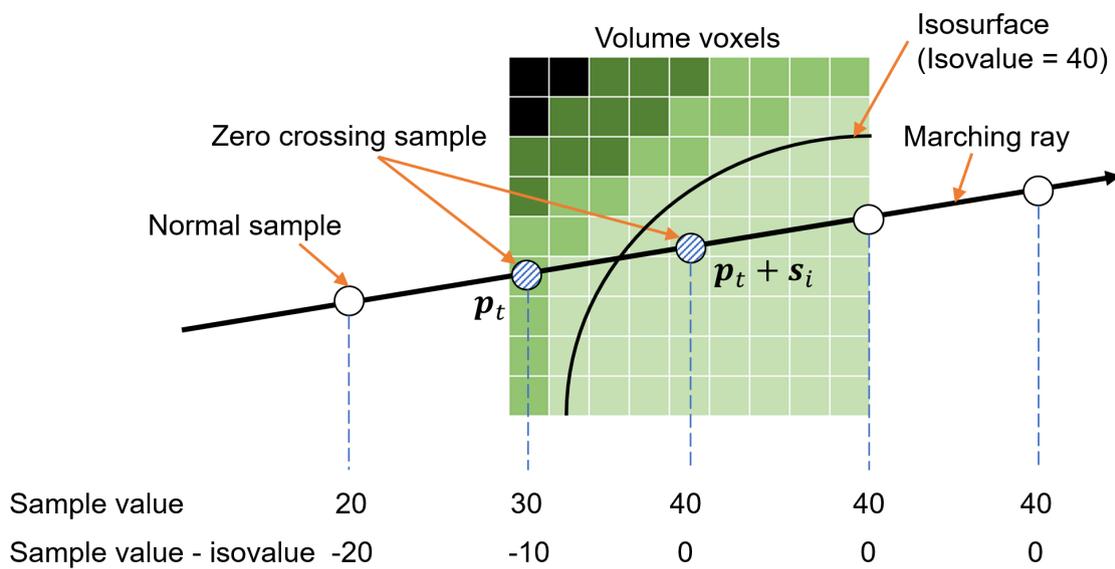


Figure B.3: An illustration of how the isosurface boundary is determined in pseudo-isosurface ray casting. Adapted from Movania [22].

B.3.3 Bisection based refinement

Bisection based refinement considers two sample points and an isovalue. Through an iterative approach, the exact location of the isovalue between these two samples is determined. At each iteration, the midpoint between the two samples is calculated and compared to the isovalue. If the value at the midpoint is less than the isovalue, it is swapped with the left sample, otherwise it is swapped with the right sample. In this way the values between the two samples can quickly be searched to find the isosurface location.

B.3.4 Calculate the gradient

Once the refined point on the isosurface boundary has been determined, the gradient, which is normal to the isosurface, is calculated at that point using the *central difference approximation*. The gradient is needed for the lighting calculations which is discussed in the next section.

$$\mathbf{N}_x = \frac{v(x - \delta, y, z) - v(x + \delta, y, z)}{2\delta} \quad (\text{B.7})$$

$$\mathbf{N}_y = \frac{v(x, y - \delta, z) - v(x, y + \delta, z)}{2\delta} \quad (\text{B.8})$$

$$\mathbf{N}_z = \frac{v(x, y, z - \delta) - v(x, y, z + \delta)}{2\delta} \quad (\text{B.9})$$

In Equations B.7-B.9, $v()$ refers to the volume data value at the given location, and x , y and z are the coordinate points of the refined boundary location. δ is a small fixed value, such as 0.005, that is used as an offset. The normal vector (\mathbf{N}) is then normalised to unit magnitude after its calculation.

B.3.5 Phong lighting

The Phong lighting model produces perceptually realistic lighting results and consists of three components, namely ambient, diffuse and specular lighting. These components, as well as the resulting image with the components combined, are shown in Figure B.4. The *ambient* component simulates the light that has been reflected and scattered around so many times that it appears to be emanating uniformly from all directions. The *diffuse* component simulates the effect of a directional light source on an object. The more the local surface of an object faces the light source, the brighter it becomes. In order to calculate this, the normal vector of each primitive surface must be known. Primitive surfaces whose normal vectors coincide more closely with the incident light will be brighter. The *specular* component simulates the reflective bright spots of light that appear on shiny surfaces, based on the light direction. This also requires the normal of each primitive surface to be known [17, 144].

Since the normal vector is computed for each fragment on the bounding box, which represent the pseudo-isosurface, it can be combined with the vectors indicated the viewing direction and the light direction to calculate the specular component's power and the diffuse component's brightness. The diffuse colour is chosen to be the same as the colour of the underlying fluorescence channel. In order to ensure that the sample is always lit appropriately, regardless of the viewing angle, we choose the light direction to be the same as the viewing direction.

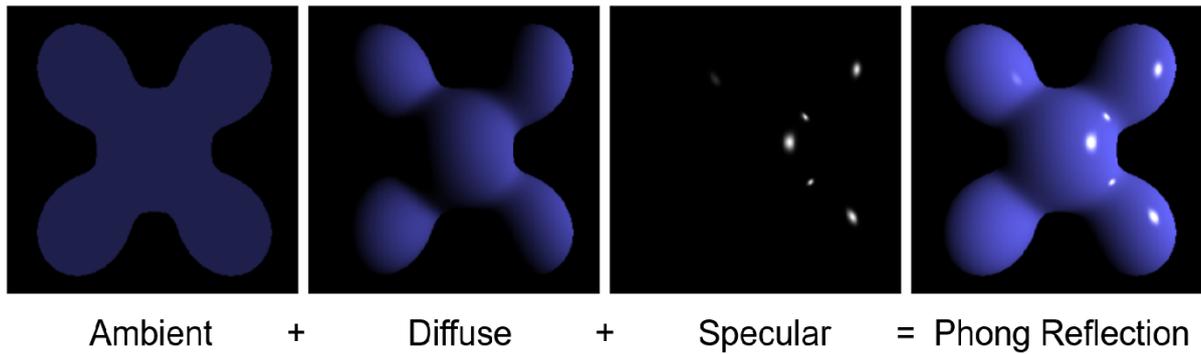


Figure B.4: An illustration of the Phong shading model. Ambient light refers to the reflected environmental light. Diffuse light refers to the direction impact of light from the light source. Sepecular light refers to the reflective brights spots as they appear on shiny objects. Adapted from [145].

B.4 Conclusion

In this appendix we discussed the implementation of the volume ray casting direct volume rendering method. We considered two variants of this method, namely standard volume ray casting as well as pseudo-isosurface volume ray casting. For each of these we showed the different calculation steps required to reconstruct the volume data three-dimensionally. In Appendix A, we consider an alternative direct volume rendering method known as texture-based volume rendering or texture slicing.

Appendix C

MEL synthetic example

In this appendix we demonstrate the steps of the MEL, as was presented in Chapter 6, algorithm by applying it to two synthetically generated image frames. This is intended only to clarify the MEL algorithm, and does not represent a realistic scenario as would be expected when analysing mitochondrial events. Although MEL is intended for application to three-dimensional samples, we will consider two-dimensional images here for clarity.

Since our images are synthetic, we omit the normalisation step shown in Figure 6.1 (in Chapter 6) and begin by showing binarised Frames 1 and 2 in Figure C.1. We also overlay these binarised images to make it easier to identify those structures that will fuse or undergo fission from Frame 1 to Frame 2. Next, each structure in the binarised frames is given a unique label number, with 0 indicating the background. Each labelled structure is then separated to create an array of images (not shown in Figure C.1) which is Gaussian filtered to allow for a less strict structure overlap matching between Frame 1 and Frame 2. The array of labelled images are also Canny filtered, leaving only the pixels on the edge which is used to determine the distance between structures and the location of the fission and fusion events. Figure C.1 shows only the first labelled structure in the array for each frame. These images are then analysed according to the process in Figure 6.2 to produce a set of locations and types of mitochondrial events that are hypothesised to have occurred in the time between Frame 1 and Frame 2. The result is then overlaid with Frame 1 using colour labels for the different events.

The matrices and arrays calculated for the synthetic example by the process depicted in Figure 6.2 (in Chapter 6) are shown in Figure C.2. The overlap matrix V is calculated by multiplying all combinations of the Gaussian filtered structures, one from Frame 1 and one from Frame 2, to determine a representation of the volume. Label number 0 indicates the background and therefore is left blank throughout.

In the synthetic example shown, there is a small overlap between Frame 1 label 4 and Frame 2 label 2 due to the blur which the filter has introduced. To compensate for such coincidental matches, all overlapping volumes that account for less than 1% of the volume of either structure in question are eliminated, and is consequently shown as 0 in the table (Figure C.2).

From matrix V , the arrays A_1 and A_2 are determined by simply reducing V to indicate which structures in the other frame presented with a non-zero overlapping volume. The relative percentage overlap, P_1 and P_2 , of each structure in one frame with all associated structures in the other frame can then be calculated from A_1 and A_2 . This is in effect a normalisation of the volumes in matrix V , where the volume of a certain structure combination is divided by the total volume of the given structure in either Frame 1 (producing P_1) or Frame 2 (producing P_2). Each row, therefore, sums to 100%.

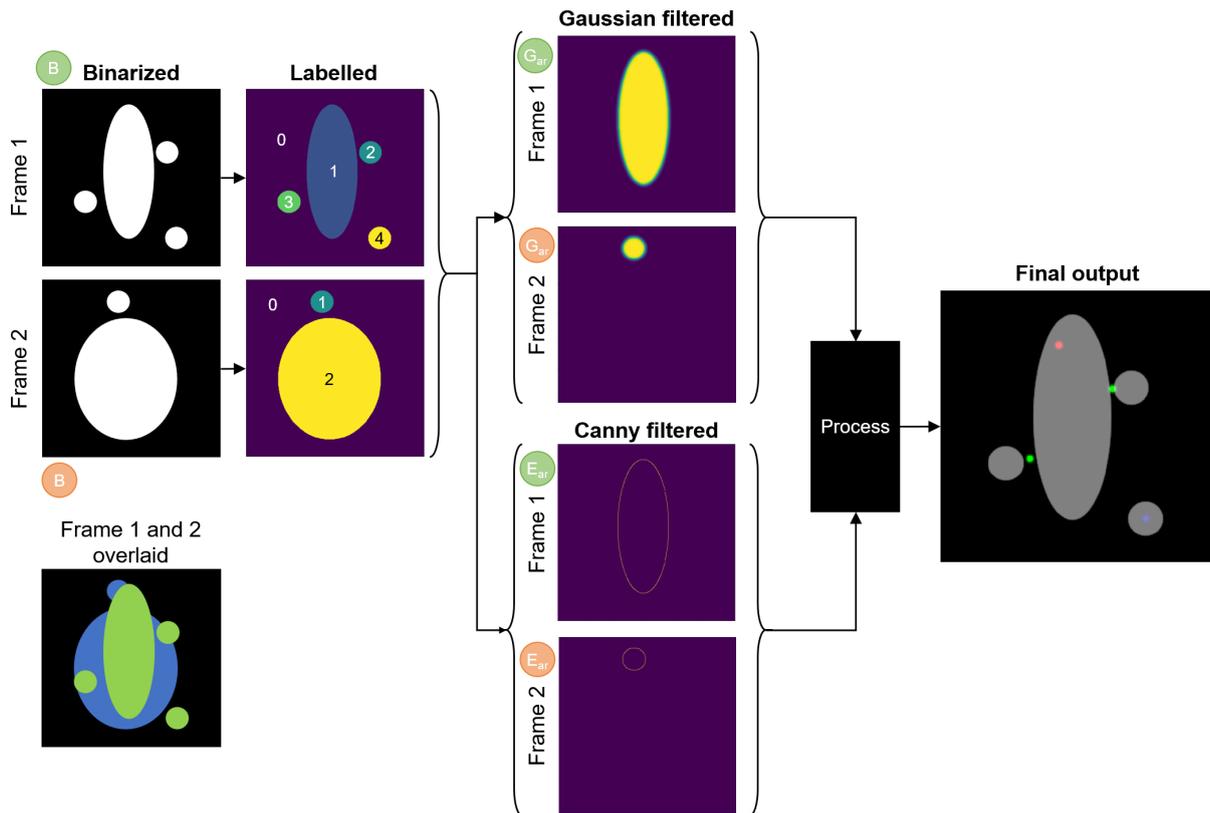


Figure C.1: Synthetic image example process flow.

Using the back-and-forth structure matching described in Section 6.2.2.2 and Figure 6.3, we generate W_1 and W_2 , to indicate which structures are associated with each other in the same frame. These are candidates for fission and fusion events, although some of them could be false matches that the next step aims to eliminate. It is worth noting that each structure combination occurs twice in the array of lists. Even though we show the calculations that follow for both, it is only necessary to use one.

Using W_1 and W_2 along with the edge array of images E_{ar} for Frame 1 and Frame 2, we find the shortest distances, D_1 and D_2 , as well as the midway points, M_1 and M_2 , between each combination of structures. If the shortest distance between the candidate structures is above a set threshold (in the case of this synthetic sample this was 50 pixels), the two structures are considered unrelated and is ignored in the visualisation (refer to Figure C.1 and Figure C.3). Secondly, if the midway point is sampled from the binarised images, B_1 and B_2 , and coincides with another structure it is also considered a false match. This is because two structures cannot fuse through a third separating structure. Rather, they would individually fuse with the central structure. A similar logic applies when considering fission events in Frame 2. This is illustrated in Figure C.3 for a fusion, fission and unrelated label combination.

Now the mitochondrial event status can be determined for each candidate structure combination in W_1 and W_2 . Structures in Frame 1 are labelled “Fuse”, “Depolarize”, or “Unrelated”, while structures in Frame 2 are labelled “Fission”, or “Unrelated”.

Finally, using the midway points, as well as the centre of mass of the structures to indicate the location of the mitochondrial event, along with the status arrays, the final output image is generated (Figure C.1).

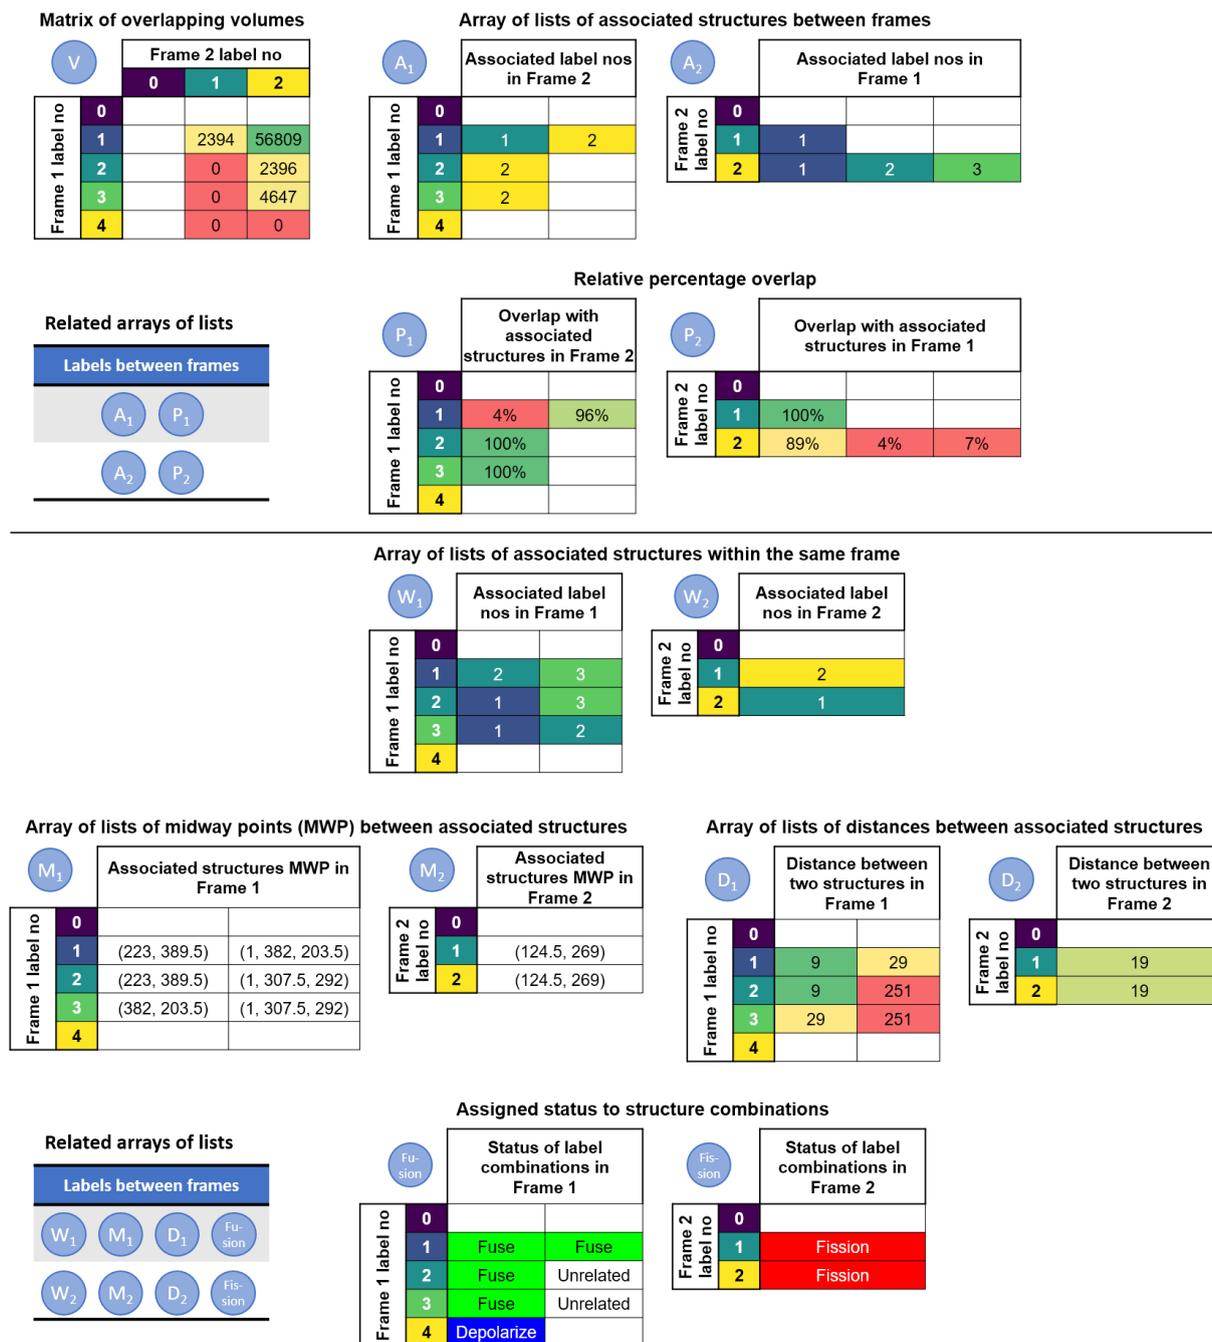


Figure C.2: Synthetic image example calculated matrix and list of arrays.

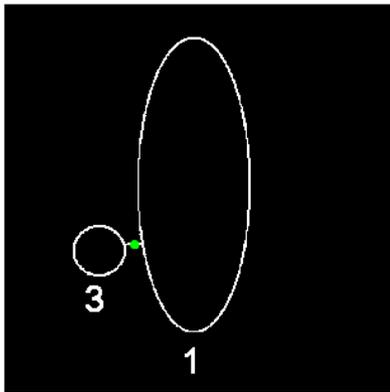
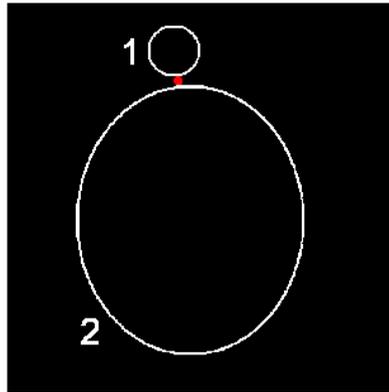
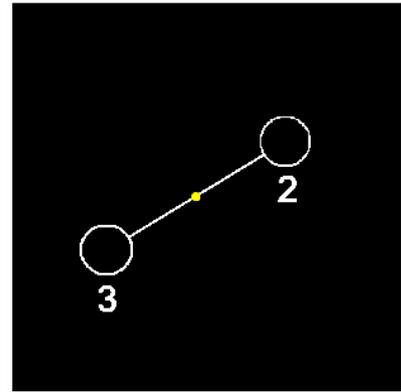
A. Frame 1
Fusion**B. Frame 2**
Fission**C. Frame 1**
Unrelated

Figure C.3: Example of determining fission and fusion event locations.

Bibliography

- [1] Wilson, T.: *Confocal microscopy*. Academic Press, 1990.
- [2] Sensen, C. and Hallgrímsson, B.: *Advanced Imaging in Biology and Medicine: Technology, Software Environments, Applications*. Springer, 2009.
- [3] Van Liere, R., De Leeuw, W., Mulder, J., Verschure, P., Visser, A., Manders, E. and Van Driel, R.: Virtual reality in biological microscopic imaging. *2002 IEEE International Symposium on Biomedical Imaging*, pp. 879–882, 2002.
- [4] Regan, M. and Pose, R.: Priority rendering with a virtual reality address recalculation pipeline. In: *SIGGRAPH '94 Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, pp. 155–162. 1994.
- [5] Vazquez-Duchêne, M., Freis, O., Denis, A., Mion, C., Jeanmaire, C., Mine, S. and Pauly, G.: Virtual reality for skin exploration. In: *VRIC '13 Proceedings of the Virtual Reality International Conference: Laval Virtual Article No. 5*. 2013.
- [6] Conte, G., Allen, Q., Forbes, A., Ajilore, O. and Leow, A.: BRAINtrinsic: A virtual reality-compatible tool for exploring intrinsic topologies of the human brain connectome. In: *International Conference on Brain Informatics and Health*, pp. 67–76. Springer, 2015.
- [7] Green, D.R., Galluzzi, L. and Kroemer, G.: Metabolic control of cell death. *Science*, vol. 345, no. 6203, p. 1250256, 2014.
- [8] Theart, R.P., Loos, B. and Niesler, T.R.: Virtual reality assisted microscopy data visualization and colocalization analysis. *BMC Bioinformatics*, vol. 18, no. 2, p. 64, 2017.
- [9] Theart, R.P., Loos, B., Powrie, Y.S. and Niesler, T.R.: Improved region of interest selection and colocalization analysis in three-dimensional fluorescence microscopy samples using virtual reality. *PLOS one*, vol. 13, no. 8, p. e0201965, 2018.
- [10] Theart, R.P., Loos, B. and Niesler, T.R.: Regression adjusted colocalisation colour mapping (racc): A novel biological visual analysis method for qualitative colocalisation analysis of 3d fluorescence micrographs. *PLOS one*, vol. 14, no. 11, 2019.
- [11] Lee, V.W., Kim, C., Chhugani, J., Deisher, M., Kim, D., Nguyen, A.D., Satish, N., Smelyanskiy, M., Chennupaty, S., Hammarlund, P. *et al.*: Debunking the 100x GPU vs. CPU myth: an evaluation of throughput computing on CPU and GPU. *ACM SIGARCH Computer Architecture News*, vol. 38, no. 3, pp. 451–460, 2010.
- [12] Genovese, L.: Graphic processing units: A possible answer to HPC. In: *4th ABINIT Developer Workshop*. 2009.
- [13] Wright, R.S.J., Haemel, N., Sellers, G. and Lipchak, B.: *OpenGL SuperBible*. 5th edn. Addison-Wesley, 2011.

- [14] Unity Real-Time Development Platform | 3D, 2D VR & AR Visualizations. Accessed 04 September 2019.
Available at: <https://unity.com/>
- [15] Shreiner, D., Sellers, G., Kessenich, J. and Licea-Kane, B.: *OpenGL Programming Guide*. 8th edn. Addison-Wesley, 2013.
- [16] Neider, J., Davis, T. and Woo, M.: *OpenGL programming guide*, vol. 14. Addison-Wesley Reading, MA, 1993.
- [17] De Vries, J.: *Learn OpenGL*. Joey de Vries, 2015.
- [18] Akenine-Moller, T., Haines, E. and Hoffman, N.: *Real-Time Rendering*. 3rd edn. A. K. Peters, Ltd., Natick, MA, USA, 2008. ISBN 1568814240, 9781568814247.
- [19] Rougier, N.P.: *Scientific visualisation: Python & Matplotlib*. Nicolas P. Rougier, 2018.
- [20] Engel, K., Hadwinger, M., Kniss, J.M., Rezk-Salama, C. and Weiskopf, D.: *Real-Time Volume Graphics*. 1st edn. AK Peters, Ltd., 2006.
- [21] Fernando, R.: *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*. 1st edn. Addison-Wesley Professional, 2004. Chapter 39. Volume Rendering Techniques.
- [22] Movania, M.M.: *OpenGL development cookbook*. Birmingham, United Kingdom: Packt Publishing Ltd, 2013.
- [23] Proakis, J.G. and Manolakis, D.G.: *Digital Signal Processing: Principles, Algorithms, and Applications*. 4th edn. Pearson Education International, 2007.
- [24] Levoy, M.: Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, vol. 8, no. 2, pp. 29–37, 1988.
- [25] Max, N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, vol. 1, no. 2, pp. 97–108, 1995.
- [26] Engel, K., Hadwiger, M., Kniss, J.M., Lefohn, A.E., Salama, C.R. and Weiskopf, D.: Real-time volume graphics. In: *ACM SIGGRAPH 2004 Course Notes*, SIGGRAPH '04. ACM, New York, NY, USA, 2004.
Available at: <http://doi.acm.org/10.1145/1103900.1103929>
- [27] Daniel, W.: *GPU-Based Interactive Visualization Techniques*. Springer Science & Business Media. Pearson Education International, 2006.
- [28] Thetawave: Volume ray casting. 2006. [Online; accessed August 31, 2019].
Available at: https://upload.wikimedia.org/wikipedia/commons/e/ec/Volume_ray_casting.png
- [29] Hadwiger, M., Ljung, P., Salama, C.R. and Ropinski, T.: Advanced illumination techniques for GPU-based volume raycasting. In: *ACM SIGGRAPH 2009 Courses*, SIGGRAPH '09, pp. 2:1–2:166. ACM, New York, NY, USA, 2009.
Available at: <http://doi.acm.org/10.1145/1667239.1667241>
- [30] Wittenbrink, C.M., Malzbender, T. and Gross, M.E.: Opacityweighted color interpolation for volume sampling. In: *Proceedings of IEEE Symposium on Volume Visualization*, pp. 135–142. 1998.

- [31] Lorensen, W.E. and Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics*, vol. 21, no. 4, pp. 163–169, July 1987.
- [32] Chu, A., Fu, C.-W. and Hanson, A.J.: GL4D: A GPU-based Architecture for Interactive 4D Visualization. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1587–1594, November/December 2009.
- [33] Cabral, B., Cam, N. and Foran, J.: Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In: *Proc. of IEEE Symposium on Volume Visualization*, pp. 91–98. 1994.
- [34] Fesharaki, H., Rezaei, L., Farrahi, F., Banihashem, T. and Jahanbakhshi, A.: Normal interpupillary distance values in an iranian population. *Journal of Ophthalmic & Vision Research*, vol. 7, no. 3, p. 231, 2012.
- [35] Kennedy, R.S., Lane, N.E., Berbaum, K.S. and Lilienthal, M.G.: Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology*, vol. 3, no. 3, pp. 203–220, 1993.
- [36] Kleciak, M.: Seamless textures. 2012.
Available at: http://wiki.openstreetmap.org/wiki/User:Marek_kleciak/Seamless_Textures
- [37] Dunn, K., Kamocka, M. and McDonald, J.: A practical guide to evaluating colocalization in biological microscopy. *Am. J. Physiol. Cell Physiol*, vol. 300, no. 4, pp. C723–C742, 2011.
- [38] Pearson, K.: Vii. mathematical contributions to the theory of evolution.-iii. regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character*, , no. 187, pp. 253–318, 1896.
- [39] Rasband, W.: *ImageJ*. National Institutes of Health, Bethesda, Maryland, USA, 1997–2004. Accessed 04 March 2016.
Available at: <http://rsb.info.nih.gov/ij/>
- [40] Oculus. Accessed 07 March 2016.
Available at: <https://www.oculus.com/en-us/>
- [41] Leap Motion | Mac & PC Motion Controller for Games, Design, Virtual Reality & More. Accessed 07 March 2016.
Available at: <https://www.leapmotion.com/>
- [42] Febretti, A., Nishimoto, A., Thigpen, T., Talandis, J., Long, L., Pirtle, J., Peterka, T., Verlo, A., Brown, M., Plepys, D., Sandin, D., Renambot, L., Johnson, A. and Leigh, J.: CAVE2: a hybrid reality environment for immersive simulation and information analysis. In: *PIE 8649, The Engineering Reality of Virtual Reality 2013, 864903*. March 2013.
- [43] Forbes, A., Villegas, J., Almryde, K. and Plante, E.: A stereoscopic system for viewing the temporal evolution of brain activity clusters in response to linguistic stimuli. In: *IS&T/SPIE Electronic Imaging*, pp. 90110I–90110I. International Society for Optics and Photonics, 2014.
- [44] Kip: Pinch Move | Leap Motion Developers. May 2016. Accessed 06 July 2016.
Available at: <https://developer.leapmotion.com/gallery/pinch-move>

- [45] Kip: UI Input Module | Leap Motion Developers. June 2016. Accessed 06 July 2016. Available at: <https://developer.leapmotion.com/gallery/ui-input-module>
- [46] Motion, L.: VR Best Practices Guidelines. June 2015. Accessed 06 July 2016. Available at: <https://developer.leapmotion.com/vr-best-practices>
- [47] Kip: Blocks | Leap Motion Developers. February 2016. Accessed 17 July 2016. Available at: <https://developer.leapmotion.com/gallery/blocks>
- [48] Spiegel, M. and Stephens, L.: *Theory and problems of statistics*. 3rd edn. McGraw-Hill, 1999.
- [49] Yates, D., Moore, D. and McCabe, G.: *The practice of statistics*. 1st edn. W.H. Freeman, 1999.
- [50] Du Toit, J.: *Statistiese metodes*. Kosmo Uitgewery Edms Bpk, 1989.
- [51] Box, G.: Non-normality and test on variances. *Biometrika*, vol. 40, no. 3/4, pp. 318–335, 1953.
- [52] Giepmans, B.N., Adams, S.R., Ellisman, M.H. and Tsien, R.Y.: The fluorescent toolbox for assessing protein location and function. *Science*, vol. 312, no. 5771, pp. 217–224, 2006.
- [53] Hood, L., Heath, J.R., Phelps, M.E. and Lin, B.: Systems biology and new technologies enable predictive and preventative medicine. *Science*, vol. 306, no. 5696, pp. 640–643, 2004.
- [54] van der Voort, H., Brakenhoff, G. and Baarslag, M.: Three-dimensional visualization methods for confocal microscopy. *Journal of Microscopy*, vol. 153, no. 2, pp. 123–132, 1989.
- [55] Van der Voort, H., Messerli, J., Noordmans, H. and Smeulders, A.: Volume visualization for interactive microscopic image analysis. *Bioimaging*, vol. 1, no. 1, pp. 20–29, 1993.
- [56] Dan, P., Cheung, J.C., Scriven, D.R. and Moore, E.D.: Epitope-dependent localization of estrogen receptor α , but not- β , in en face arterial endothelium. *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 284, no. 4, pp. H1295–H1306, 2003.
- [57] Scriven, D.R., Klimek, A., Asghari, P., Bellve, K. and Moore, E.D.: Caveolin-3 is adjacent to a group of extradyadic ryanodine receptors. *Biophysical Journal*, vol. 89, no. 3, pp. 1893–1901, 2005.
- [58] Fletcher, P.A., Scriven, D.R., Schulson, M.N. and Moore, E.D.: Multi-image colocalization and its statistical significance. *Biophysical Journal*, vol. 99, no. 6, pp. 1996–2005, 2010.
- [59] Lifshitz, L.M., Collins, J.A., Moore, E.D. and Gauch, J.: Computer vision and graphics in fluorescence microscopy. In: *Proceedings of IEEE Workshop on Biomedical Image Analysis*, pp. 166–175. IEEE, 1994.
- [60] Wu, Y., Eghbali, M., Ou, J., Lu, R., Toro, L. and Stefani, E.: Quantitative determination of spatial protein-protein correlations in fluorescence confocal microscopy. *Biophysical Journal*, vol. 98, no. 3, pp. 493–504, 2010.
- [61] Comeau, J.W., Kolin, D.L. and Wiseman, P.W.: Accurate measurements of protein interactions in cells via improved spatial image cross-correlation spectroscopy. *Molecular BioSystems*, vol. 4, no. 6, pp. 672–685, 2008.

- [62] Lumkwana, D., Du Toit, A., Kinnear, C. and Loos, B.: Autophagic flux control in neurodegeneration: Progress and precision targeting - where do we stand? *Progress in Neurobiology*, vol. 153, pp. 64–85, 2017.
- [63] Cataldo, A.M., Peterhoff, C.M., Troncoso, J.C., Gomez-Isla, T., Hyman, B.T. and Nixon, R.A.: Endocytic pathway abnormalities precede amyloid β deposition in sporadic Alzheimer's disease and down syndrome: differential effects of APOE genotype and presenilin mutations. *The American Journal of Pathology*, vol. 157, no. 1, pp. 277–286, 2000.
- [64] Nixon, R.A., Wegiel, J., Kumar, A., Yu, W.H., Peterhoff, C., Cataldo, A. and Cuervo, A.M.: Extensive involvement of autophagy in alzheimer disease: an immuno-electron microscopy study. *Journal of Neuropathology & Experimental Neurology*, vol. 64, no. 2, pp. 113–122, 2005.
- [65] Nixon, R.A. and Yang, D.-S.: Autophagy failure in Alzheimer's disease-locating the primary defect. *Neurobiology of Disease*, vol. 43, no. 1, pp. 38–45, 2011.
- [66] Lee, S., Sato, Y. and Nixon, R.A.: Primary lysosomal dysfunction causes cargo-specific deficits of axonal transport leading to alzheimer-like neuritic dystrophy. *Autophagy*, vol. 7, no. 12, pp. 1562–1563, 2011.
- [67] Torres, M., Jimenez, S., Sanchez-Varo, R., Navarro, V., Trujillo-Estrada, L., Sanchez-Mejias, E., Carmona, I., Davila, J.C., Vizuet, M., Gutierrez, A. *et al.*: Defective lysosomal proteolysis and axonal transport are early pathogenic events that worsen with age leading to increased APP metabolism and synaptic Abeta in transgenic APP/PS1 hippocampus. *Molecular Neurodegeneration*, vol. 7, no. 1, p. 59, 2012.
- [68] Avila, J.: Microtubule functions. *Life Sciences*, vol. 50, no. 5, pp. 327–334, 1992.
- [69] Jahreiss, L., Menzies, F.M. and Rubinsztein, D.C.: The itinerary of autophagosomes: from peripheral formation to kiss-and-run fusion with lysosomes. *Traffic*, vol. 9, no. 4, pp. 574–587, 2008.
- [70] Heintzmann, R. and Cremer, C.G.: Laterally modulated excitation microscopy: improvement of resolution by using a diffraction grating. In: *Optical Biopsies and Microscopic Techniques III*, vol. 3568, pp. 185–196. International Society for Optics and Photonics, 1999.
- [71] Manders, E., Stap, J., Brakenhoff, G., Van Driel, R. and Aten, J.: Dynamics of three-dimensional replication patterns during the s-phase, analysed by double labelling of DNA and confocal microscopy. *Journal of Cell Science*, vol. 103, no. 3, pp. 857–862, 1992.
- [72] Demandolx, D. and Davoust, J.: Multicolour analysis and local image correlation in confocal microscopy. *Journal of Microscopy*, vol. 185, no. 1, pp. 21–36, 1997.
- [73] Manders, E., Verbeek, F. and Aten, J.: Measurement of co-localization of objects in dual-colour confocal images. *Journal of Microscopy*, vol. 169, no. 3, pp. 375–382, 1993.
- [74] Dunn, K.W., Kamocka, M.M. and McDonald, J.H.: A practical guide to evaluating colocalization in biological microscopy. *American Journal of Physiology-Cell Physiology*, vol. 300, no. 4, pp. C723–C742, 2011.
- [75] Fox, M., Arndt-Jovin, D., Jovin, T., Baumann, P. and Robert-Nicoud, M.: Spatial and temporal distribution of dna replication sites localized by immunofluorescence and confocal microscopy in mouse fibroblasts. *Journal of Cell Science*, vol. 99, no. 2, pp. 247–253, 1991.

- [76] Dutartre, H., Davoust, J., Gorvel, J.-P. and Chavrier, P.: Cytokinesis arrest and redistribution of actin-cytoskeleton regulatory components in cells expressing the Rho GTPase CDC42Hs. *Journal of Cell Science*, vol. 109, no. 2, pp. 367–377, 1996.
- [77] Mietelska-Porowska, A., Wasik, U., Goras, M., Filipek, A. and Niewiadomska, G.: Tau protein modifications and interactions: their role in function and dysfunction. *International Journal of Molecular Sciences*, vol. 15, no. 3, pp. 4671–4713, 2014.
- [78] Hamano, T., Gendron, T.F., Causevic, E., Yen, S.-H., Lin, W.-L., Isidoro, C., DeTure, M. and Ko, L.-w.: Autophagic-lysosomal perturbation enhances tau aggregation in transfectants with induced wild-type tau expression. *European Journal of Neuroscience*, vol. 27, no. 5, pp. 1119–1130, 2008.
- [79] Rodríguez-Martín, T., Cuchillo-Ibáñez, I., Noble, W., Nyenya, F., Anderton, B.H. and Hanger, D.P.: Tau phosphorylation affects its axonal transport and degradation. *Neurobiology of Aging*, vol. 34, no. 9, pp. 2146–2157, 2013.
- [80] Lee, S., Kim, W., Li, Z. and Hall, G.F.: Accumulation of vesicle-associated human tau in distal dendrites drives degeneration and tau secretion in an in situ cellular tauopathy model. *International Journal of Alzheimer's Disease*, vol. 2012, 2012.
- [81] Neethling, A., Engelbrecht, L., Loos, B., Kinnear, C., Theart, R., Abrahams, S., Niesler, T., Mellick, G.D., Williams, M. and Bardiën, S.: Wild-type and mutant (g2019s) leucine-rich repeat kinase 2 (lrrk2) associate with subunits of the translocase of outer mitochondrial membrane (tom) complex. *Experimental Cell Research*, vol. 375, no. 2, pp. 72–79, 2019.
- [82] Zinchuk, V., Zinchuk, O. and Okada, T.: Quantitative colocalization analysis of multi-color confocal immunofluorescence microscopy images: pushing pixels to explore biological phenomena. *Acta Histochemica et Cytochemica*, pp. 0707300003–0707300003, 2007.
- [83] Bolte, S. and Cordelières, F.P.: A guided tour into subcellular colocalization analysis in light microscopy. *Journal of Microscopy*, vol. 224, no. 3, pp. 213–232, 2006.
- [84] Costes, S.V., Daelemans, D., Cho, E.H., Dobbin, Z., Pavlakis, G. and Lockett, S.: Automatic and quantitative measurement of protein-protein colocalization in live cells. *Biophysical Journal*, vol. 86, no. 6, pp. 3993–4003, 2004.
- [85] Agnati, L.F., Fuxe, K., Torvinen, M., Genedani, S., Franco, R., Watson, S., Nussdorfer, G.G., Leo, G. and Guidolin, D.: New methods to evaluate colocalization of fluorophores in immunocytochemical preparations as exemplified by a study on A2A and D2 receptors in chinese hamster ovary cells. *Journal of Histochemistry & Cytochemistry*, vol. 53, no. 8, pp. 941–953, 2005.
- [86] Villalta, J.I., Galli, S., Iacaruso, M.F., Arciuch, V.G.A., Poderoso, J.J., Jares-Erijman, E.A. and Pietrasanta, L.I.: New algorithm to determine true colocalization in combination with image restoration and time-lapse confocal microscopy to MAP kinases in mitochondria. *PLOS one*, vol. 6, no. 4, p. e19031, 2011.
- [87] Jaskolski, F., Mülle, C. and Manzoni, O.J.: An automated method to quantify and visualize colocalized fluorescent signals. *Journal of Neuroscience Methods*, vol. 146, no. 1, pp. 42–49, 2005.
- [88] Gorlewicz, A., Włodarczyk, J., Wilczek, E., Gawlak, M., Cabaj, A., Majczynski, H., Nestorowicz, K., Herbik, M.A., Grieb, P., Slawinska, U. *et al.*: CD44 is expressed in non-myelinating schwann cells of the adult rat, and may play a role in neurodegeneration-induced glial plasticity at the neuromuscular junction. *Neurobiology of Disease*, vol. 34, no. 2, pp. 245–258, 2009.

- [89] Ostler, N., Britzen-Laurent, N., Liebl, A., Naschberger, E., Lochnit, G., Ostler, M., Forster, F., Kunzelmann, P., Ince, S., Supper, V. *et al.*: Gamma interferon-induced guanylate binding protein 1 is a novel actin cytoskeleton remodeling factor. *Molecular and Cellular Biology*, vol. 34, no. 2, pp. 196–209, 2014.
- [90] Bettgazzi, B., Bellani, S., Roncon, P., Guarnieri, F.C., Bertero, A., Codazzi, F., Valtorta, F., Simonato, M., Grohovaz, F. and Zacchetti, D.: eIF4B phosphorylation at ser504 links synaptic activity with protein translation in physiology and pathology. *Scientific Reports*, vol. 7, no. 1, p. 10563, 2017.
- [91] Lemieux, M., Labrecque, S., Tardif, C., Labrie-Dion, É., LeBel, É. and De Koninck, P.: Translocation of CaMKII to dendritic microtubules supports the plasticity of local synapses. *J Cell Biol*, vol. 198, no. 6, pp. 1055–1073, 2012.
- [92] De Juan-Sanz, J., Núñez, E., Villarejo-López, L., Pérez-Hernández, D., Rodríguez-Fraticelli, A.E., López-Corcuera, B., Vázquez, J. and Aragón, C.: Na⁺/K⁺-ATPase is a new interacting partner for the neuronal glycine transporter GlyT2 that downregulates its expression in vitro and in vivo. *Journal of Neuroscience*, vol. 33, no. 35, pp. 14269–14281, 2013.
- [93] Cornbleet, P.J. and Gochman, N.: Incorrect least-squares regression coefficients in method-comparison analysis. *Clinical Chemistry*, vol. 25, no. 3, pp. 432–438, 1979.
- [94] Deming, W.E.: Statistical adjustment of data. 1943.
- [95] Adcock, R.J.: A problem in least squares. *The Analyst*, vol. 5, no. 1, pp. 53–54, 1878.
- [96] Cheng, C.-L., Van Ness, J.W. *et al.*: *Statistical regression with measurement error*. London: Arnold and New York: Oxford University Press, 1999.
- [97] Hammond, G.R., Takasuga, S., Sasaki, T. and Balla, T.: The ML1N_{x2} phosphatidylinositol 3, 5-bisphosphate probe shows poor selectivity in cells. *PLOS one*, vol. 10, no. 10, p. e0139957, 2015.
- [98] Christoforides, C., Rainero, E., Brown, K.K., Norman, J.C. and Toker, A.: PKD controls $\alpha\text{v}\beta\text{3}$ integrin recycling and tumor cell invasive migration through its substrate rabaptin-5. *Developmental Cell*, vol. 23, no. 3, pp. 560–572, 2012.
- [99] Du Toit, A., Hofmeyr, J.-H.S., Gniadek, T.J. and Loos, B.: Measuring autophagosome flux. *Autophagy*, vol. 14, no. 6, pp. 1060–1071, 2018.
- [100] Loos, B., Engelbrecht, A.-M., Lockshin, R.A., Klionsky, D.J. and Zakeri, Z.: The variability of autophagy and cell death susceptibility: Unanswered questions. *Autophagy*, vol. 9, no. 9, pp. 1270–1285, 2013.
- [101] Du Toit, A., De Wet, S., Hofmeyr, J.-H., Müller-Nedebock, K. and Loos, B.: The precision control of autophagic flux and vesicle dynamics—a micropattern approach. *Cells*, vol. 7, no. 8, p. 94, 2018.
- [102] McCarron, J.G., Wilson, C., Sandison, M.E., Olson, M.L., Girkin, J.M., Saunter, C. and Chalmers, S.: From structure to function: mitochondrial morphology, motion and shaping in vascular smooth muscle. *Journal of Vascular Research*, vol. 50, no. 5, pp. 357–371, 2013.
- [103] Twig, G., Hyde, B. and Shirihai, O.S.: Mitochondrial fusion, fission and autophagy as a quality control axis: the bioenergetic view. *Biochimica et Biophysica Acta (BBA)-Bioenergetics*, vol. 1777, no. 9, pp. 1092–1097, 2008.

- [104] Gilkerson, R.W., De Vries, R.L., Lebot, P., Wikstrom, J.D., Torgykes, E., Shirihai, O.S., Przedborski, S. and Schon, E.A.: Mitochondrial autophagy in cells with mtdna mutations results from synergistic loss of transmembrane potential and mtorc1 inhibition. *Human Molecular Genetics*, vol. 21, no. 5, pp. 978–990, 2011.
- [105] Mishra, P. and Chan, D.C.: Metabolic regulation of mitochondrial dynamics. *J Cell Biol*, vol. 212, no. 4, pp. 379–387, 2016.
- [106] Vidoni, S., Zanna, C., Rugolo, M., Sarzi, E. and Lenaers, G.: Why mitochondria must fuse to maintain their genome integrity. *Antioxidants & Redox Signaling*, vol. 19, no. 4, pp. 379–388, 2013.
- [107] Chen, H., Vermulst, M., Wang, Y.E., Chomyn, A., Prolla, T.A., McCaffery, J.M. and Chan, D.C.: Mitochondrial fusion is required for mtdna stability in skeletal muscle and tolerance of mtdna mutations. *Cell*, vol. 141, no. 2, pp. 280–289, 2010.
- [108] Twig, G., Elorza, A., Molina, A.J., Mohamed, H., Wikstrom, J.D., Walzer, G., Stiles, L., Haigh, S.E., Katz, S., Las, G. *et al.*: Fission and selective fusion govern mitochondrial segregation and elimination by autophagy. *The EMBO Journal*, vol. 27, no. 2, pp. 433–446, 2008.
- [109] Legros, F., Lombès, A., Frachon, P. and Rojo, M.: Mitochondrial fusion in human cells is efficient, requires the inner membrane potential, and is mediated by mitofusins. *Molecular Biology of the Cell*, vol. 13, no. 12, pp. 4343–4354, 2002.
- [110] Twig, G., Liu, X., Liesa, M., Wikstrom, J.D., Molina, A.J., Las, G., Yaniv, G., Hajnóczky, G. and Shirihai, O.S.: Biophysical properties of mitochondrial fusion events in pancreatic β -cells and cardiac cells unravel potential control mechanisms of its selectivity. *American Journal of Physiology-Cell Physiology*, vol. 299, no. 2, pp. C477–C487, 2010.
- [111] Narendra, D., Tanaka, A., Suen, D.-F. and Youle, R.J.: Parkin is recruited selectively to impaired mitochondria and promotes their autophagy. *The Journal of Cell Biology*, vol. 183, no. 5, pp. 795–803, 2008.
- [112] Mortiboys, H., Thomas, K.J., Koopman, W.J., Klaffke, S., Abou-Sleiman, P., Olpin, S., Wood, N.W., Willems, P.H., Smeitink, J.A., Cookson, M.R. *et al.*: Mitochondrial function and morphology are impaired in parkin-mutant fibroblasts. *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, vol. 64, no. 5, pp. 555–565, 2008.
- [113] Benard, G., Bellance, N., Jose, C., Melser, S., Nouette-Gaulain, K. and Rossignol, R.: Multi-site control and regulation of mitochondrial energy production. *Biochimica et Biophysica Acta (BBA)-Bioenergetics*, vol. 1797, no. 6-7, pp. 698–709, 2010.
- [114] Berman, S., Pineda, F.J. and Hardwick, J.M.: Mitochondrial fission and fusion dynamics: the long and short of it. *Cell Death and Differentiation*, vol. 15, no. 7, p. 1147, 2008.
- [115] Karbowski, M., Arnoult, D., Chen, H., Chan, D.C., Smith, C.L. and Youle, R.J.: Quantitation of mitochondrial dynamics by photolabeling of individual organelles shows that mitochondrial fusion is blocked during the bax activation phase of apoptosis. *J Cell Biol*, vol. 164, no. 4, pp. 493–499, 2004.
- [116] Kriel, J., Müller-Nedebock, K., Maarman, G., Mbizana, S., Ojuka, E., Klumperman, B. and Loos, B.: Coordinated autophagy modulation overcomes glioblastoma chemoresistance through disruption of mitochondrial bioenergetics. *Scientific Reports*, vol. 8, no. 1, p. 10348, 2018.

- [117] Westrate, L.M., Drocco, J.A., Martin, K.R., Hlavacek, W.S. and MacKeigan, J.P.: Mitochondrial morphological features are associated with fission and fusion events. *PLOS one*, vol. 9, no. 4, p. e95265, 2014.
- [118] Rambold, A.S., Kostecky, B., Elia, N. and Lippincott-Schwartz, J.: Tubular network formation protects mitochondria from autophagosomal degradation during nutrient starvation. *Proceedings of the National Academy of Sciences*, vol. 108, no. 25, pp. 10190–10195, 2011.
- [119] Wang, S., Xiao, W., Shan, S., Jiang, C., Chen, M., Zhang, Y., Lü, S., Chen, J., Zhang, C., Chen, Q. *et al.*: Multi-patterned dynamics of mitochondrial fission and fusion in a living cell. *PLOS one*, vol. 7, no. 5, p. e19879, 2012.
- [120] Peng, J.-Y., Lin, C.-C., Chen, Y.-J., Kao, L.-S., Liu, Y.-C., Chou, C.-C., Huang, Y.-H., Chang, F.-R., Wu, Y.-C., Tsai, Y.-S. *et al.*: Automatic morphological subtyping reveals new roles of caspases in mitochondrial dynamics. *PLOS Computational Biology*, vol. 7, no. 10, p. e1002212, 2011.
- [121] Cagalinec, M., Safiulina, D., Liiv, M., Liiv, J., Choubey, V., Wareski, P., Veksler, V. and Kaasik, A.: Principles of the mitochondrial fusion and fission cycle in neurons. *J Cell Sci*, vol. 126, no. 10, pp. 2187–2197, 2013.
- [122] Valente, A.J., Maddalena, L.A., Robb, E.L., Moradi, F. and Stuart, J.A.: A simple imagej macro tool for analyzing mitochondrial network morphology in mammalian cell culture. *Acta Histochemica*, vol. 119, no. 3, pp. 315–326, 2017.
- [123] Leonard, A.P., Cameron, R.B., Speiser, J.L., Wolf, B.J., Peterson, Y.K., Schnellmann, R.G., Beeson, C.C. and Rohrer, B.: Quantitative analysis of mitochondrial morphology and membrane potential in living cells using high-content imaging, machine learning, and morphological binning. *Biochimica et Biophysica Acta (BBA)-Molecular Cell Research*, vol. 1853, no. 2, pp. 348–360, 2015.
- [124] Reis, Y., Bernardo-Faura, M., Richter, D., Wolf, T., Brors, B., Hamacher-Brady, A., Eils, R. and Brady, N.R.: Multi-parametric analysis and modeling of relationships between mitochondrial morphology and apoptosis. *PLOS one*, vol. 7, no. 1, p. e28694, 2012.
- [125] Sarder, P. and Nehorai, A.: Deconvolution methods for 3-D fluorescence microscopy images. *IEEE Signal Processing Magazine*, vol. 23, no. 3, pp. 32–45, 2006.
- [126] Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [127] Canny, J.: A computational approach to edge detection, readings in computer vision: issues, problems, principles, and paradigms. *IEEE Trans. Pattern*, 1987.
- [128] Ntsapi, C., Swart, C., Lumkwana, D. and Loos, B.: Autophagic flux failure in neurodegeneration: Identifying the defect and compensating flux offset. *Autophagy in Current Trends in Cellular Physiology and Pathology*, p. 157, 2016.
- [129] Maday, S. and Holzbaur, E.L.: Compartment-specific regulation of autophagy in primary neurons. *Journal of Neuroscience*, vol. 36, no. 22, pp. 5933–5945, 2016.
- [130] Imamura, H., Nhat, K.P.H., Togawa, H., Saito, K., Iino, R., Kato-Yamada, Y., Nagai, T. and Noji, H.: Visualization of atp levels inside single living cells with fluorescence resonance energy transfer-based genetically encoded indicators. *Proceedings of the National Academy of Sciences*, vol. 106, no. 37, pp. 15651–15656, 2009.

- [131] Sheng, Z. and Cai, Q.: Mitochondrial transport in neurons: impact on synaptic homeostasis and neurodegeneration. *Nature Reviews Neuroscience*, vol. 13, no. 2, p. 77, 2012.
- [132] Knott, A.B., Perkins, G., Schwarzenbacher, R. and Bossy-Wetzell, E.: Mitochondrial fragmentation in neurodegeneration. *Nature Reviews Neuroscience*, vol. 9, no. 7, p. 505, 2008.
- [133] Jin, S.M. and Youle, R.J.: Pink1-and parkin-mediated mitophagy at a glance. *J Cell Sci*, vol. 125, no. 4, pp. 795–799, 2012.
- [134] Diot, A., Hinks-Roberts, A., Lodge, T., Liao, C., Dombi, E., Morten, K., Brady, S., Fratter, C., Carver, J., Muir, R. *et al.*: A novel quantitative assay of mitophagy: Combining high content fluorescence microscopy and mitochondrial dna load to quantify mitophagy and identify novel pharmacological tools against pathogenic heteroplasmic mtdna. *Pharmacological Research*, vol. 100, pp. 24–35, 2015.
- [135] Liu, X., Weaver, D., Shirihai, O. and Hajnóczky, G.: Mitochondrial 'kiss-and-run': interplay between mitochondrial motility and fusion-fission dynamics. *The EMBO Journal*, vol. 28, no. 20, pp. 3074–3089, 2009.
- [136] Tyas, L., Brophy, V.A., Pope, A., Rivett, A.J. and Tavaré, J.M.: Rapid caspase-3 activation during apoptosis revealed using fluorescence-resonance energy transfer. *EMBO Reports*, vol. 1, no. 3, pp. 266–270, 2000.
- [137] Galluzzi, L., Zamzami, N., Rouge, T.d.L.M., Lemaire, C., Brenner, C. and Kroemer, G.: Methods for the assessment of mitochondrial membrane permeabilization in apoptosis. *Apoptosis*, vol. 12, no. 5, pp. 803–813, 2007.
- [138] Nikolaisen, J., Nilsson, L.I., Pettersen, I.K., Willems, P.H., Lorens, J.B., Koopman, W.J. and Tronstad, K.J.: Automated quantification and integrative analysis of 2d and 3d mitochondrial shape and network properties. *PLOS one*, vol. 9, no. 7, p. e101365, 2014.
- [139] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems*, pp. 2672–2680. 2014.
- [140] Mirza, M. and Osindero, S.: Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [141] Isola, P., Zhu, J.-Y., Zhou, T. and Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1125–1134. 2017.
- [142] Sarkar, S.: Chemical screening platforms for autophagy drug discovery to identify therapeutic candidates for huntington's disease and other neurodegenerative disorders. *Drug Discovery Today: Technologies*, vol. 10, no. 1, pp. e137–e144, 2013.
- [143] Joseph, J., Seervi, M., Sobhan, P.K. and Retnabai, S.T.: High throughput ratio imaging to profile caspase activity: potential application in multiparameter high content apoptosis analysis and drug screening. *PLOS one*, vol. 6, no. 5, p. e20114, 2011.
- [144] Wolff, D.: *OpenGL 4.0 Shading Language Cookbook*. Birmingham, United Kingdom: Packt Publishing Ltd, 2011.
- [145] Smith, B: Phong reflection model. 2006. [Online; accessed November 12, 2019]. Available at: https://en.wikipedia.org/wiki/Phong_reflection_model#/media/File:Phong_components_version_4.png